Chapter 1
Affine Permutations and an Affine Coxeter Monoid

Tom Denton

Abstract In this expository paper, we describe results on pattern avoidance arising from the affine Catalan monoid. The schema of affine codes as canonical decompositions in conjunction with two-row moves is detailed, and then applied in studying the Catalan quotient of the 0-Hecke monoid. We prove a conjecture of Hanusa and Jones concerning periodicity in the number of fully-commutative affine permutations. We then re-frame prior results on fully commutative elements using the affine codes.

1.1 Introduction

The Hecke Algebra of a Weyl group is a deformation by a parameter $q = \frac{q_1}{q_2}$ which for generic values yields an algebra with representation theory equivalent to that of the group algebra of the original Weyl group. Our interest is ultimately in the affine symmetric group, $\tilde{S}_n$, which has index set $I = \{0, 1, 2, \ldots, n-1\} = \mathbb{Z}_n$, but we will describe the basic constructions of algebras and monoids in full generality, and then specialize where necessary.

At $q_1 = 0$, however, we obtain the 0-*Hecke algebra*, which can be interpreted as a monoid algebra of the 0-*Hecke monoid*. If the original Weyl group is generated by some simple reflections s_i for i int he index set I, then the 0-Hecke monoid is generated by π_i for $i \in I$. The commutation and braid relations between the pi_i match the relations on the s_i, but we have $\pi_i^2 = \pi_i$ instead of $s_i^2 = 1$. Thus, the 0-Hecke monoid is generated by projections, instead of reflections. There is a bijection between elements of the orginal group and elements of the monoid, but the representation theory changes considerably. This representation theory was initially studied by Norton for the 0-Hecke algebra of the symmetric group [Nor79], and expanded to arbitrary finite Weyl type by Carter [Car86]. Later, the 0-Hecke algebra was shown to have characters determined by the pairing of noncommutative and quasi-symmetric functions [KT97], and was explored as a special case of the representation theory of $\mathscr{J}$-trivial monoids [DHST11].

At $q_1 = q_2 = 0$, one obtains the *nilHecke algebra*, which, like the 0-Hecke algebra, can be considered as a monoid algebra. This *NilHecke monoid* is generated by a_i for $i \in I$, again with the same commutation and braid relations as the original group, but now with $a_i^2 = 0$. The nilHecke algebra has proven useful in studying reduced word combinatorics, since any non-reduced word in the a_i evaluates to 0. It has also proven important in the categorification of quantum algebras [KL09][Rou08], and provides a very useful model for the study of k-Schur functions [Lam06][LLM$^+$12][LM03].

In studying the product formula on the k-Schur functions, the present author developed a new combinatorial model for the affine symmetric group [Den12], combining aspects of the inversion vector or affine code and RC-graphs, originally developed to study reduced word combinatorics [BB93]. The model provides an interpretation of the inversion vector as a unique maximal decomposition of the affine permutation into cyclically decreasing elements, originally introduced in [Lam06]. This combinatorial model is equally functional in the affine symmetric group, as well as its 0-Hecke monoid and nilHecke monoid. This model was immediately used for calculating a special case of the k-Littlewood-Richardson rule.

In the author's dissertation, a certain quotient of the 0-Hecke algebra for the symmetric group was studied in relation to pattern avoidance. This led to an algebraic interpretation of certain kinds of pattern avoidance via the fibers of the quotient map. The quotient is known as the Catalan monoid, since it has Catalan-many elements, and each fiber of the quotient contains a unique maximal-length element avoiding the pattern $[2,3,1]$, and a unique minimal-length element avoiding the pattern $[3,2,1]$. The author also extended this result to the affine setting, defining an affine Catalan monoid, and generalizing the $[3,2,1]$ avoidance result to that setting. This provides a bijection between $[2,3,1]$ and $[3,2,1]$-avoiding permutations. The bijection is equivalent to the bijection of Simion and Schmidt [SS85], a fundamental early result in pattern avoidance, but we place the bijection in a new algebraic setting.

In the present article, we will review the connection between cyclic decompositions of affine permutations and the affine code, the connection between the 0-Hecke algebra and pattern avoidance, and link the two topics via 'shadow diagrams.' Along the way, we will confirm a conjecture of Hanusa and Jones[HJ11] concerning periodicity in the length generating function of the fully commutative elements of the affine symmetric group. This is the content of Theorem 2.

1.1.1 Acknowledgements

The new results in this paper mainly developed from conversations with Brant Jones. I am always grateful for the existence of Sage **??** and Sage-Combinat **??**, which I use extensively in computer exploration. I am also grateful to York University and the

Fields Institute for providing space and support in 2011-12, and Maseno University, Kenya, and the US Department of State for providing support in 2012-13.

1.2 Background and Notation

1.2.1 Affine Permutations, Hecke Algebras, Specializations

An n-affine permutation (we usually omit the n when no ambiguity will arise) is a permutation $x : \mathbb{Z} \to \mathbb{Z}$ satisfying the following properties:

- $\sum_{i=1}^{n} x(i) = \binom{n+1}{2}$, and
- $x(i+n) = x(i) + n$.

One can show easily that such permutations form a group, which is called the *affine symmetric group* $\tilde{S}_n$.

The second condition implies that one may completely specify an affine permutation simply by specifying $x(i)$ for i in the set $\{1,2,\dots,n\}$. The list $[x(1),\dots,x(n)]$ is called the *window notation* for x. One can observe that a list of n integers is a valid window notation if and only if the sum is $\binom{n+1}{2}$, and when each entry of the list is reduced mod n one obtains a permutation of $\mathbb{Z}_n$. (If any residues were repeated, then the second condition could be used to show that the presumed permutation has repeated entries, and is thus not a permutation.) For example, $[5,-2,3]$ is a valid window notation for an affine permutation (the sum is 6, and the list reduces to $[2,1,0]$ modulo 3). On the other hand, $[6,-3,3]$ is not a window for an affine permutation σ. The sum is 6, but we would have:

$$\sigma(8) = \sigma(2+3+3) = \sigma(2)+3+3 = -3+3+3 = 3 = \sigma(3),$$

so that σ is not a bijection. We will usually identify an affine permutation with its window notation.

Generators and relations for $\tilde{S}_n$ are given as follows. There is one simple reflection s_i for each $i \in \mathbb{Z}_n$, being the permutation which exchanges $mn+i$ and $mn+i+1$ for every $m \in \mathbb{Z}$, while fixing all other numbers. The action of the simple reflection is either on the left or the right: The left action exchanges the values, while the right action exchanges the numbers in the given positions. For example, if we take $x = [1,3,2]$, then $s_1 x = [2,3,1]$, exchanging the values 1 and 2. On the other hand, $x s_1 = [3,1,2]$, exchanging the first and second positions.

These generators satisfy the following relations:

- $s_i^2 = 1$,
- $s_i s_j = s_j s_i$ for all $j > i$ with $j - i > 1$, and
- $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ for all i.

Note that by omitting the generator s_0, we recover a group isomorphic to the usual symmetric group.

The Iwahori-Hecke algebra $H_q(\tilde{S}_n)$ is a q-deformation of the group algebra $\mathbb{C}\tilde{S}_n$, generated by elements T_i for $i \in \mathbb{Z}_n$ with relations:

- $T_i^2 = (q-1)T_i + q$,
- $T_iT_j = T_jT_i$ for all $j > i$ with $j - i > 1$, and
- $T_iT_{i+1}T_i = T_{i+1}T_iT_{i+1}$ for all i.

In short, the quadratic relation is deformed while the relations between different generators are preserved. At $q = 1$, we recover $\mathbb{C}\tilde{S}_n$. At $q = 0$, we have $T_i^2 = -T_i$. For convenience, we set $\pi_i := -T_i$, so that $\pi_i^2 = \pi_i$. It is easy to see that the π_i satisfy the commutation and braid relations. The monoid generated by the π_i is then called the 0-Hecke monoid $\tilde{H}_n$. There is a bijection between elements of this monoid and the set of affine permutations. In particular, reduced words in the affine symmetric group are also reduced in $\tilde{H}_n$.

The π_i may be considered as *anti-sorting* operators on the collection of affine permutations, transposing values if they are in order, and leaving them fixed if not. Thus, if $x = [3,1,2]$, then $\pi_1 x = [3,2,1]$, anti-sorting the values 1 and 2. On the other hand, $x\pi_1 = [3,1,2] = x$, since the first and second positions are already anti-sorted.

For expedience, we will often write words in the generators as subscripts on the generator. Furthermore, we will avoid examples with $n > 10$ in this paper, and will thus may omit commas in the writing of lists of indices. Thus, $\pi_1\pi_2\pi_1$ may be written as π_{121} unambiguously.

1.2.2 Pattern Avoidance

A permutation may be thought of as a sequence of numbers, and an affine permutation may be thought of as a doubly-infinite sequence of numbers. Let $\sigma = [\sigma_1, \dots, \sigma_k]$ be a permutation, and x a permutation or affine permutation. We say that x *contains* σ if there exist $i_1 < i_2 < \cdots < i_k$ such that $x_{i_1}, x_{i_1}, \dots, x_{i_k}$ are in the same relative order as σ. If x does not contain σ, then we say that x **avoids** σ, or that x is **σ-avoiding.**

For example, the pattern $[1,2]$ appears in any x such that there exists a $x_i < x_j$ for some $i < j$. The only $[12]$-avoiding permutation in S_N, then, is the long element, which is strictly decreasing in one-line notation. As a larger example, the permutation $[\mathbf{3},\mathbf{4},5,\mathbf{2},1,6]$ contains the pattern $[231]$ at the bold positions. In fact, this permutation contains six distinct instances of the pattern $[231]$.

Of particular interest is the pattern $[3,2,1]$. A permutation (or affine permutation) which avoids $[3,2,1]$ is called *fully commutative*, or FC for short. The name arises because one can show that if x is FC, then any reduced word for x can be obtained from any other via a sequence of commutation relations [BJS93] [Lam06] [Gre02].

1.3 Canonical Decompositions of Affine Permutations

Given a subset $A \subsetneq \mathbb{Z}_n$ with $|A| = m < n$ we define the *cyclically decreasing element* d_A (d for 'down') to be the product $d_A := T_{i_1} \cdots T_{i_n}$ for $i_l \in A$, where if $j, j-1 \in A$ then j appears to the left of $j-1$ in any reduced word for d_A. (One may similarly define cyclically increasing elements, where j appears to the right of $j-1$.) Note that such products may be specialized to either the affine symmetric group or the 0-Hecke monoid.

For example, if we take $n = 9$ and $A = \{0,1,3,4,5,8\} \subset \mathbb{Z}_9$ at $q = 0$, we have $d_A = \pi_{108543}$. The idea here is that we want to write the product of the generators in a decreasing list, but the the index set $\mathbb{Z}_n$ is circular, having no greatest element. So we rely on a local notion of order between adjacent elements. There is a bijection between cyclically decreasing elements d_A and subsets of $\mathbb{Z}_n$. The cyclically decreasing elements are also known as *Pieri factors* in the literature, and play an important role in the theory of k-Schur functions.

Every affine permutation can be written as a product of cyclically decreasing elements, $x = d_{A_l} d_{A_{l-1}} \cdots d_{A_1}$. Such a decomposition is then called an *α-decomposition of x*. We write $\mathbf{A} = (A_1, A_2, \ldots, A_l)$, and may then form a composition $\mathrm{sh}((A)) = (|A_1|, \ldots, |A_n|)$. A given x may have many different α-decompositions. (In fact, the affine Stanley symmetric functions are defined as a sum over all such decompositions [Lam06].) we may order the α-decompositions lexicographically (or reverse lexicographically) by comparison of $\mathrm{sh}((A))$. Under this ordering, we have the following theorem:

Theorem 1 (Canonical Cyclically Decreasing Decomposition.[Lam06][Den12]). *Every affine permutation x admits a unique maximal decomposition under the anti-lexicographic ordering as a product of cyclically decreasing elements $x = d_{\mathbf{A}}$. This decomposition has $A_{i+1} \subset \{j-1 \mid j \in A_i\}$ for each i, and thus $\mathrm{sh}(\mathbf{A})$ is a partition.*

We've made two choices so far: a choice between cyclically increasing and decreasing elements in building the decomposition of x, and the choice of whether to find the maximal decomposition according to lexicographic or anti-lexicographic ordering on $\mathrm{sh}(\mathbf{A})$. We'll henceforth call the maximal lexicographic decomposition the *left* decomposition, and the anti-lexicographic decomposition the *right* decomposition. Thus, the theorem is stated for the right decreasing decomposition, and may be modified for any of the other three choices. (In particular, the containment property for the A_i must be modified for other cases, though an analogous statement holds.)

The maximal decomposition is closely related to the *affine code* of the permutation, which is also known as the inversion vector or affine Lehmer code. And there are actually four different affine codes one can associate to any affine permutation, corresponding to the four possible choices. The right decreasing code $\mathrm{RD}(x)$ is the list $[c_1, c_2, \ldots, c_n]$ where c_i is the number of $j < i$ with $x(j) > x(i)$. One can show that the numbers c_i are always finite, and that one of the c_i must be equal to 0.

One may recover the maximal decomposition for an affine permutation x from the affine code by carrying out the following steps:

- Make a Ferrer's diagram of the code of x. This diagram has n columns, and in each column has c_i boxes.
- Fill each box with a *residue*, by filling the jth box from the bottom in the c_i column with the number $i-j \mod n$. (We start counting j from 1, so the bottommost residue in the c_i column is just $i-1$; we then count backwards up the column modulo n.)
- Now each *row* of the resulting diagram corresponds to a cyclically decreasing element, obtained by reading cyclically right-to-left starting from any empty column. (One should imagine the diagram drawn on a cylinder, just as the cyclically decreasing elements come from the 'circular' index set $\mathbb{Z}_n$.)

Example 1. Let $n=5$, and x be the affine permutation with window notation $[4,-4,7,5,3]$. Then the affine code of x is $[0,6,0,1,3]$. For example, there are six elements to the left of the entry $x(2)=-4$ which are larger than -4. Specifically, these are $x(-7)=-3, x(-4)=-1, x(-2)=2, x(-1)=0$, and $x(1)=4$. The diagram of the permutation is pictured in Figure 1.1. (The grey areas are the 'shadows' of the columns, described later.) Starting with one of the empty columns, we read right-to-left, top-to-bottom, considering the diagram as though it were on a cylinder. The maximal decomposition of x is then:

$$x = d_{\{1\}}d_{\{2\}}d_{\{3\}}d_{\{2,4\}}d_{\{3,0\}}d_{\{4,3,1\}}.$$

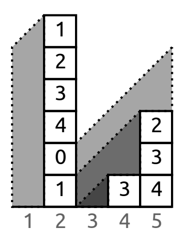

Fig. 1.1 Diagram of the maximal decomposition of x with window notation $[4,-4,7,5,3]$.

By using a very small modification of the 'moves' on RC-diagrams considered by Bergeron and Billey [BB93], one can move between different α-decompositions of x. There are two kinds of 'moves' available, by which one can move boxes in the diagram into other rows. A *commutation move* moves a box with no neighbours above or to the right up into the next row directly, using only the commutation relations. A *chute move* moves a box past a two-by-l block of boxes, changing the residue of the given box, using a sequence of braid relations. (The precise rule for this is obtained by repeated application of the basic braid relations in ASn_n.) This is illustrated in 1.2, with an example in 1.3.

How can we apply this schema to affine pattern avoidance? It is known that an affine permutation is $[3,2,1]$-avoiding if and only if it has no reduced words in which

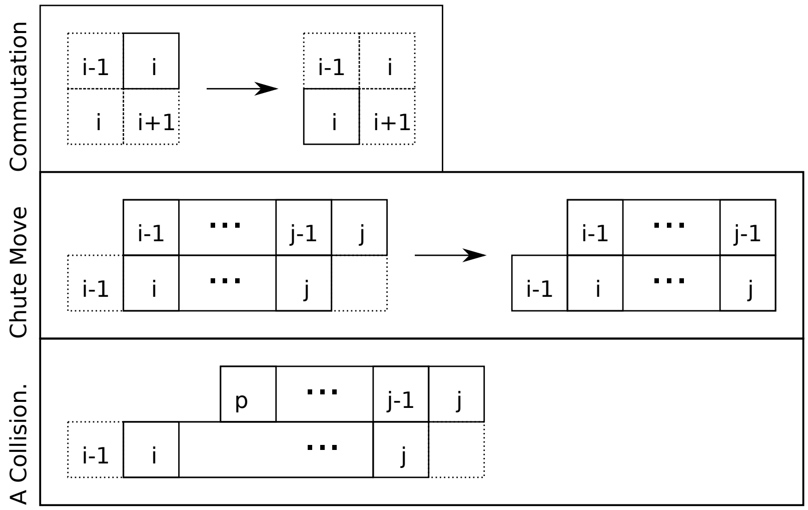

Fig. 1.2 The two-row moves on an α-decomposition. The collision shows a situation where a chute move will lead to a situation involving squaring a generator. This doesn't happen in reduced decompositions, and the result depends on whether one is working in the affine symmetric group (in which case the two colliding boxes annihilate one another), the nilHecke monoid (in which case the whole permutation is equal to 0), or the 0-Hecke monoid (where one of the boxes is removed while the other remains), according to the various specializations of T_i^2 at different values of q.

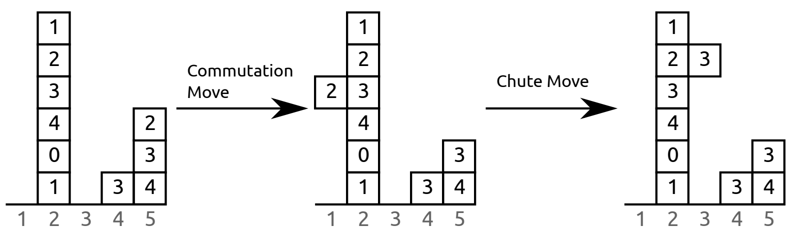

Fig. 1.3 Examples of other α-decompositions of x with window notation $[4,-4,7,5,3]$. The leftmost picture is the maximal decomposition of x. The middle is obtained from the maximal α-decomposition by a commutation move. And the rightmost is obtained from the middle by a chute move (in this case, a simple braid relation).

one may apply a braid relation. This was shown independently by Lam [Lam06] and Green [Gre02]. Given a diagram of a maximal decomposition, we can convert it to a *shadow diagram*. This simply involves drawing the shadows of each column, extending down-to-the-left at 45 degrees. The shadows also wrap-around the diagram. (Refer back to Figure 1.1 for an example.) One can then observe the following:

Proposition 1. *An affine permutation x is $[3,2,1]$-avoiding if and only if no column of the diagram of x is completely in the shadow of another column. Equivalently, $c_{i-j} >= c_i - j$ for all $i, j \in \mathbb{Z}_n$, considering the indices as elements of $\mathbb{Z}_n$.*

Thus, the example permutation with window $[4,-4,7,5,3]$ contains a $[3,2,1]$ pattern, since $c_5 = c_{2-2} = 3 < c_2 - 2 = 4$. This allows for a very fast check of whether a given affine permutation is $[3,2,1]$ avoiding. It takes linear time to construct the affine code (indeed, only one pass through the main window is required), and only one pass through the code to determine whether the permutation is FC.

Proof. If any column is completely shadowed by another column, then the top box in the shadowed column can be brought into position for a braid move. Thus, a permutation with a completely shadowed column is not fully commutative.

For the converse, we show that if x is not fully commutative, then x has a shadowed column. Every α-decomposition of x can be obtained from the maximal decomposition by a sequence of two-row moves [Den12], which includes all reduced words for x. Let $w = [w_1, w_2, \dots, w_l]$ be a reduced word for x with a consecutive subword $s_i s_{i+1} s_i$ occurring farthest to the left amongst all reduced words for x; in particular, let the braid $s_i s_{i+1} s_i$ occur in positions w_j, w_{j+1}, w_{j+2}, with j minimal amongst all reduced words for x. One may then find the maximal decomposition of the permutation with word $[w_{j+3}, \dots, w_l]$. Then inserting the three letters w_j, w_{j+1}, w_{j+2} will ensure that the column with top box w_j is fully shadowed by the column with top box w_{j+2}. The column which now has top box w_j cannot grow any larger through the addition of further boxes from $[w_1, \dots, w_{j-1}]$ without the application of a braid move, which is impossible by the minimality of j. Therefore, there is a completely shadowed column in the maximal decomposition of x, as desired.

Corollary 1. *Each FC affine permutation has diagram given by a union of partitions, separated by empty columns, such that shadow condition is satisfied.*

Proof. If $0 \neq c_i < c_{i+1}$, then the shadow condition is violated. Thus, for every i, we have $c_i > c_{i+1}$ or $c_i = 0$.

1.3.1 Enumeration of Fully Commutative Affine Permutations.

We can also use this structure to study the length-enumeration of the fully commutative elements. It was shown by Crites that there are infinitely many $[3,2,1]$-avoiding elements [Cri10], but one can still ask how many affine permutations there are of length l. Let $F_n(l)$ be the number of fully commutative affine permutations of length l in $\tilde{S}_n$. This was studied by Jones and Hanusa via generating functions [HJ11]. In particular, they were able to show that the number of FC elements of length l is eventually periodic in l. Observationally, this periodic behaviour began at $l = 1 + \lceil \frac{n-1}{2} \rceil \lfloor \frac{n-1}{2} \rfloor$. We now show that this bound is sharp.

Suppose that x has code c. Then let $U(x)$ be the affine permutation with code obtained by adding one to each non-zero entry of c. Likewise, let $D(x)$ be the permutation with code obtained by reducing each non-zero entry of c by one. For our running example with window $[4,-4,7,5,3]$, and affine code $[0,6,0,1,3]$, then $U(x)$ has code $[0,7,0,2,3]$. Thus, $U(x)$ is the permutation with window $[6,-5,8,4,2]$. Likewise, $D(x)$ has code $[0,5,0,0,2]$, which gives the permutation $[3,-3,5,6,4]$.

It is clear that if x is fully commutative then so is $U(x)$ and $D(x)$; reducing or increasing the size of each column does not affect the shadow condition of Proposition 1. It is also clear that $U(D(x)) = x$ unless the code of x has some entry c_i equal to 1. We will call these *shift-minimal* elements.

Definition 1. A fully commutative affine permutation x is *shift-minimal* if the code c of x has at least one entry equal to 1. For any fully commutative affine permutation, we set $\mathrm{col}(x)$ to be the number of non-zero entries of c, and $r(x)$ to be the length of x reduced modulo $\mathrm{col}(x)$. Set $M_{i,j}$ to be the number of shift-minimal fully commutative affine permutations x with $\mathrm{col}(x) = i$ and $r(x) = j$.

For example, let y be the affine permutation with window $[-3,-1,8,1,10]$. Then the code of y is $[4,3,0,3,0]$, and y is FC. We have $\mathrm{col}(y) = 3$, and $r(y) = 10 \mod 3 = 1$.

Lemma 1. *There are only finitely many shift-minimal FC elements. The maximal length of a shift-minimal element is* $\leq 1 + \lceil \frac{n-1}{2} \rceil \lfloor \frac{n-1}{2} \rfloor$*, and there exist elements attaining this length.*

Proof. Since each shift-minimal element contains a 1 in its code c, there can be (by the shadow condition) no $c_i > n$. This establishes an upper bound of $1+(n-2)(n-1)$ on the length of a shift-minimal element, which in turn implies that there are finitely many such elements.

To construct an element of maximal length, suppose without loss of generality that $c_1 = 1$. Then we claim an optimal strategy to construct a long element is to make a $j \times (n-j)$ rectangle to the right of the first column. (See figure 1.4.)

By Corollary 1, the diagram must be a union of partitions, $P_1, P_2, \ldots, P_l$, placed so that they satisfy the shadow rule. Since we are maximizing the total number of boxes, we may take each partition P_i and complete it to a rectangle of height equal to the first column of P_i. This operation will never violate the shadow rule. Thus, we consider a sequence of rectangles. The last rectangle may as well be pushed as far to the right as possible, so that it is adjacent (on the cylinder) to the single box in the first column.

Now consider the second-to-last rectangle. On inspection, we see that it has a maximal width k determined by the height of the last rectangle and its placement relative to the first column. (See Figure 1.5). If it has less than this maximal width, then we can add more columns to the rectangle until the maximal width is reached. Then this rectangle has width k, and some height a which is less than or equal to the height of the last rectangle. But to maximize the size of the rectangle, we should then maximize the height a. Then a is equal to the height of the last rectangle, which shows that they are, in fact, a single big rectangle. Applying this argument to all of the partitions, we see that we only need to consider a single rectangle!

Now we should then maximize the area of our single rectangle: Thus, we maximize the product $j(n-j)$ over j. The maximum product of two integers that sum to n is $\lceil \frac{n}{2} \rceil \lfloor \frac{n}{2} \rfloor$. Including the initial column of size 1, we obtain a maximal length shift-minimal element of length $1 + \lceil \frac{n}{2} \rceil \lfloor \frac{n}{2} \rfloor$, as desired.

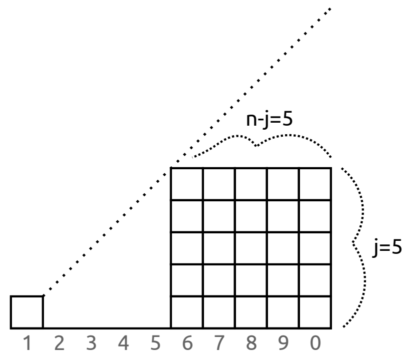

Fig. 1.4 A shift-minimal element of maximal length for $n = 10$. The single-box column ensures the element is shift-minimal. Then building a large rectangle allows us to maximize the total number of boxes. We should then maximize the size of rectangles over the possible heights j and widths $n - j$.

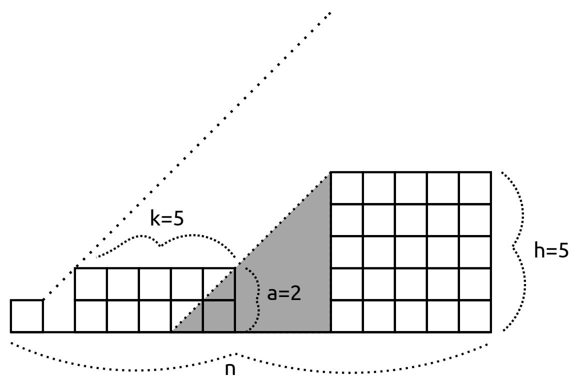

Fig. 1.5 Ruling out the two-rectangle case. The rectangle on the right is freely chosen, so as not to shadow the single box int he first column. The middle rectangle, in order to maximize the number of boxes, should have its upper-right point on the shadow cast by the right-most rectangle. Maximizing the width gives the constant k. Then the rectangle has area ak. Maximizing a then 'merges' the rectangle with the original rectangle.

Now we consider our function $F_n(l)$. It is clear that for every FC affine permutation x of length l, there exists some minimal k such that $D^k(x)$ is shift-minimal. This involves subtracting some multiple of $\mathrm{col}(x)$ until we are left with a shift-minimal element. Let $M_{i,j}$ be the number of shift-minimal FC affine permutations with $\mathrm{col}(x) = i$ and $r(x) = j$. Then we observe that for l large:

$$F_n(l) = \sum_{i=1}^{n-1} M_{i,l\%i},$$

where $l\%i$ denotes $l \mod i$. Then every $F_n(l)$ is a sum of $M_{i,j}$'s; therefore, once all of the shift-minimal elements are accounted for, the function $F_n(l)$ must begin being periodic. And this occurs exactly at $1+\lceil \frac{n}{2} \rceil \lfloor \frac{n}{2} \rfloor$. Thus we have shown:

Theorem 2. *The function $F_n(l)$ is periodic in l, with periodic behaviour beginning at $l = 1+\lceil \frac{n}{2} \rceil \lfloor \frac{n}{2} \rfloor$.*

Ostensibly, from this construction, the period of $F_n(l)$ would be $\leq n!$. In fact, it was shown by Hanusa and Jones that the period divides n. This implies that there should be some interesting relations between the numbers $M_{i,j}$ which remain to be fully explored.

1.4 Affine Catalan Monoid and Pattern Avoidance

In this section, we mainly summarize results from [Den13] and the author's thesis. We detail

An interesting quotient of the 0-Hecke monoid may be obtained by introducing the relation, for every $i \in I$:

$$\pi_{i+1}\pi_i\pi_{i+1} = \pi_i\pi_{i+1}.$$

In the finite case, this gives a monoid isomorphic to the *Catalan monoid* of functions $f : \{1,\ldots,n\} \to \{1,\ldots,n\}$ satisfying:

- $f(p) \leq p$, and
- $p \leq \Rightarrow f(p) \leq f(q)$.

These functions form a monoid under composition, called the *non-decreasing parking functions*. (See e.g. [Sol96]; it also is described under the notation $\mathscr{C}_N$ in e.g. [Pin10, Chapter XI.4] and, together with many variants, in [GM09, Chapter 14]). Similarly, it is a natural quotient of Kiselman's monoid [GM10, KM09]. In [DHST11], this monoid was studied as an instance of the larger class of order-preserving regressive functions on monoids, and a set of explicit orthogonal idempotents in the algebra was described.

The fibers of this quotient have a nice property:

Theorem 3. *Each fiber of the quotient map $H_0(S_N) \to \mathrm{NDPF}_N$ contains a unique $[3,2,1]$-avoiding element for minimal length and a unique $[2,3,1]$-avoiding element of maximal length.*

One can check that this directly implements the Simone-Schmidt bijection, which is one of the essential early results in the study of pattern avoidance. The bijection's proof was originally combinatorial, but this theorem shows that there is actually an algebraic reason for the bijection.

We also note that by taking the quotient of the 0-Hecke monoid by the relation $\pi_{i+1}\pi_i\pi_{i+1} = \pi_{i+1}\pi_i$, one obtains another monoid isomorphic to the Catalan monoid, though the fibers contain a unique $[3,1,2]$-avoiding permutation instead of a $[2,3,1]$-avoiding permutation.

We can use the affine codes of the finite permutations to directly construct the fibers. The finite permutation group is a parabolic subgroup of the affine permutation group, so all of the technology for working with affine codes descends to the finite case. Here, the code becomes a familiar Lehmer code or inversion vector, and the two-row moves become very close to moves on RC-graphs.

For finite permutations, the code is calculated in exactly the same way, but there are no π_0 generators and no 0-residues in the diagram of a finite permutation. Thus, we have $c_i \leq i-1$ for each i. Given the shadow diagram of a finite permutation, we obtain the $[3,2,1]$-avoiding permutation by deleting all columns that are completely shadowed. Likewise, to obtain the maximal, $[3,1,2]$-avoiding element, we insert fully shadowed columns. Such a fiber is pictured in Figure 1.6.

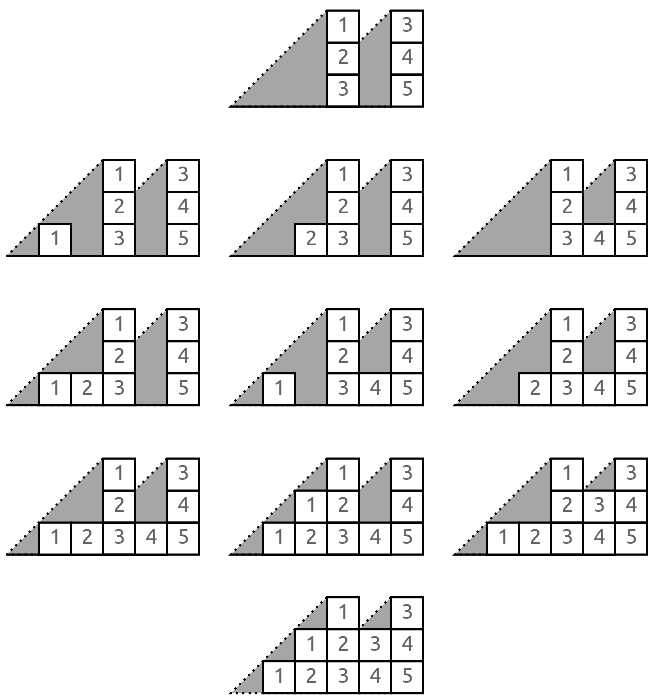

Fig. 1.6 A fiber of the finite quotient from $H_0(S_N)$ to the Catalan monoid, NDPF_N. At the top is the diagram of the $[3,2,1]$-avoiding permutation $[2,4,5,1,6,3]$. At the bottom is the $[2,3,1]$-avoiding permutation $[6,5,2,1,4,3]$. All other elements of the fiber lie between. Interestingly, we may notice that the fiber has an implicit poset structure arising from containment of diagrams.

1.4.1 Affine pattern avoidance

In the affine case, the introduction of the extra relation gives a monoid isomorphic to the affine non-decreasing parking functions.

Definition 2. The extended affine non-decreasing parking functions are the functions $f : \mathbb{Z} \to \mathbb{Z}$ which are:

- Regressive: $f(i) \leq i$,
- Order Preserving: $i \leq j \Rightarrow f(i) \leq f(j)$, and
- Skew Periodic: $f(i+N) = f(i)+N$.

Define the **shift functions** sh_t as the functions sending $i \to i-t$ for every i.

The **affine non-decreasing parking functions** $\mathrm{NDPF}^{(1)}{}_N$ are obtained from the extended affine non-decreasing parking functions by removing the shift functions for all $t \neq 0$.

It is worth noting that the quotient seems to work well for the 0-Hecke monoid, but isn't useful in the case of general q.

We can summarize the correspondence in the following theorem:

Theorem 4. *The affine non-decreasing parking functions* $\mathrm{NDPF}^{(1)}{}_N$ *are a* $\mathscr{J}$*-trivial monoid which can be obtained as a quotient of the* 0*-Hecke monoid of the affine symmetric group by the relations* $\pi_j\pi_{j+1}\pi_j = \pi_j\pi_{j+1}$*, where the subscripts are interpreted modulo* N*. Each fiber of this quotient contains a unique* $[3,2,1]$*-avoiding affine permutation.*

By a result of Crites, there are infinitely many affine permutations that avoid a pattern σ if and only if σ contains the pattern $[3,2,1]$ [Cri10]. Thus, there are infinitely many $[3,2,1]$-avoiding affine permutations, but only finitely many $[2,3,1]$-avoiding affine permutations. As in the finite case, the $[3,2,1]$-avoiding element is obtained by deleting all fully-shadowed columns. But filling in shadowed columns may be ambiguous, since there must be at least one empty column; thus we cannot expect that there is a unique element of maximal length in the fibers of the affine Catalan quotient.

References

[BB93] Nantel Bergeron and Sara Billey, *Rc-graphs and schubert polynomials*, Experiment. Math. **2** (1993), no. 4, 257–269.

[BJS93] Sara C. Billey, William Jockusch, and Richard P. Stanley, *Some combinatorial properties of Schubert polynomials*, Journal of Algebraic Combinatorics **2** (1993), 345–374, 10.1023/A:1022419800503.

[Car86] R. W. Carter, *Representation theory of the* 0*-Hecke algebra*, J. Algebra **104** (1986), no. 1, 89–103. MR MR865891 (88a:20014)

[Cri10] Andrew Crites, *Enumerating pattern avoidance for affine permutations*, Electronic Journal of Combinatorics **17** (2010), P13.

[Den12] Tom Denton, *Canonical decompositions of affine permutations, affine codes, and split k-schur functions*, the electronic journal of combinatorics **19** (2012), no. 4, P19.

[Den13] ———, *Algebraic and affine pattern avoidance*, Preprint: http://arxiv.org/abs/1303.3767.

[DHST11] Tom Denton, Florent Hivert, Anne Schilling, and Nicolas Thiéry, *The representation theory of $\mathscr{J}$-trivial monoids*, Séminaire Lotharingien de Combinatoire **64** (2011), B64d.

[GM09] Olexandr Ganyushkin and Volodymyr Mazorchuk, *Classical finite transformation semigroups*, Algebra and Applications, vol. 9, Springer-Verlag London Ltd., London, 2009. MR 2460611 (2009i:20123)

[GM10] ———, *On Kiselman quotients of the* 0*-Hecke algebra*, 16 pages, preprint arXiv:1006.0316.

[Gre02] R.M. Green, *On 321-avoiding permutations in affine weyl groups*, Journal of Algebraic Combinatorics **15** (2002), 241–252, 10.1023/A:1015012524524.

[HJ11] Christopher RH Hanusa and Brant C Jones, *The enumeration of fully commutative affine permutations*, DMTCS Proceedings (2011), no. 01, 457–468.

[KL09] Mikhail Khovanov and Aaron D Lauda, *A diagrammatic approach to categorification of quantum groups i*, Represent. Theory **13** (2009), 309–347.

[KM09] Ganna Kudryavtseva and Volodymyr Mazorchuk, *On Kiselman's semigroup*, Yokohama Math. J. **55** (2009), no. 1, 21–46. MR 2561084

[KT97] Daniel Krob and Jean-Yves Thibon, *Noncommutative symmetric functions. IV. Quantum linear groups and Hecke algebras at $q = 0$*, J. Algebraic Combin. **6** (1997), no. 4, 339–376. MR MR1471894 (99c:05196)

[Lam06] Thomas Lam, *Affine stanley symmetric functions*, Amer. J. Math **128** (2006), 1553–1586.

[LLM+12] Thomas Lam, Luc Lapointe, Jennifer Morse, Anne Schilling Mark Shimozono, and Mike Zabrocki, *A primer on k-schur functions*, 2012, In preparation.

[LM03] Luc Lapointe and Jennifer Morse, *Tableaux on $k+1$-cores, reduced words for affine permutations, and k-schur expansions*, J. Combinatorial Theory, Ser. A **112** (2003), 44–81.

[Nor79] P. N. Norton, 0*-Hecke algebras*, J. Austral. Math. Soc. Ser. A **27** (1979), no. 3, 337–357. MR MR532754 (80e:16015)

[Pin10] Jean-Éric Pin, *Mathematical Foundations of Automata Theory*, 2010.

[Rou08] Raphaël Rouquier, *2-kac-moody algebras*, arXiv preprint arXiv:0812.5023 (2008).

[Sol96] Andrew Solomon, *Catalan monoids, monoids of local endomorphisms, and their presentations*, Semigroup Forum **53** (1996), no. 3, 351–368. MR 1406781 (98c:20104)

[SS85] R. Simion and F.W. Schmidt, *Restricted permutations*, European Journal of Combinatorics (1985), 383–406.