



ELSEVIER

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta


The Non-Commutative Cycle Lemma

Craig Armstrong¹, James A. Mingo², Roland Speicher^{2,3},
Jennifer C.H. Wilson¹

Queen's University, Department of Mathematics and Statistics, Jeffery Hall, Kingston, ON K7L 3N6, Canada

ARTICLE INFO

Article history:

Received 26 March 2009

Available online 22 December 2009

Keywords:

Free group

Cyclic reduction

Kesten's law

Random matrices

ABSTRACT

We present a non-commutative version of the Cycle Lemma of Dvoretzky and Motzkin that applies to free groups and use this result to solve a number of problems involving cyclic reduction in the free group. We also describe an application to random matrices, in particular the fluctuations of Kesten's Law.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Suppose n is a positive integer and $\epsilon_1, \dots, \epsilon_n$ is a string of $+1$'s and -1 's. The string is said to be *dominating* if for each $1 \leq i \leq n$ the number of $+1$'s in the initial substring $\epsilon_1, \dots, \epsilon_i$ is more than the number of -1 's in $\epsilon_1, \dots, \epsilon_i$.

Let $k := \sum_{i=1}^n \epsilon_i$ be the difference between the number of $+1$'s and -1 's in $\epsilon_1, \dots, \epsilon_n$. The Cycle Lemma asserts that if $k > 0$ then there are exactly k cyclic permutations of the string $\epsilon_1, \dots, \epsilon_n$ which are dominating. The statement dates at least to J. Bertrand [1] in 1887. Dvoretzky and Motzkin [3, Theorem 1] gave a simple and elegant proof; see also Dershowitz and Zaks [4] for a survey of recent references and applications.⁴

In this paper we prove a version of the Cycle Lemma for free groups, which when the group has only one generator reduces to the result of Dvoretzky and Motzkin. In the free group case, the non-commutative nature of the problem means that simply counting the excess of $+1$'s to -1 's cannot describe the resulting configurations. Instead, we are required to use planar diagrams.

E-mail addresses: 5ca5@queensu.ca (C. Armstrong), mingo@mast.queensu.ca (J.A. Mingo), speicher@mast.queensu.ca (R. Speicher), 4jchw@qmlink.queensu.ca (J.C.H. Wilson).

¹ Research supported by a USRA from the Natural Sciences and Engineering Research Council of Canada.

² Research supported by Discovery Grants from the Natural Sciences and Engineering Research Council of Canada.

³ Research supported by a Killam Fellowship from the Canada Council for the Arts.

⁴ The Cycle Lemma covers also generalizations to p -dominating strings; however the case $p > 1$ follows by replacing each $\epsilon = -1$ with p ϵ 's equal to -1 .

The application to random matrices involves the concept of asymptotic freeness introduced over twenty years ago by D. Voiculescu [14], who showed that the asymptotics of certain random matrix ensembles can be described using the algebra of the free group. At the end of the paper we shall give an indication of the problem on random matrix theory that led us to the Non-Commutative Cycle Lemma.

Let \mathbb{F}_N be the free group on the N generators u_1, \dots, u_N and let $w = l_1 \cdots l_n$ be a word in $u_1^{\pm 1}, \dots, u_N^{\pm 1}$. By a word we mean a string a letters which may or may not simplify. The *length* of a word is the number of letters in the string. Following usual terminology, we shall say that $w = l_1 \cdots l_k$ is reduced, or to be more precise *linearly reduced*, if for all $1 \leq i < k$, $l_i \neq l_{i+1}^{-1}$. We shall say that w is *cyclically reduced* if in addition $l_k \neq l_1^{-1}$. Equivalently, w is cyclically reduced if $w \cdot w$ is linearly reduced. We say that a word w *reduces linearly* to a word w' if w' is linearly reduced and can be obtained from w by successively removing neighboring letters which are inverses of each other. We say that w *reduces cyclically* to w' if w' is cyclically reduced and we can obtain w' from w by successive removal of cyclic neighbors which are inverses of each other (i.e., in that case we might also remove the first and the last letter if they are inverses of each other). We say that a word w is *reducible to 1* if it reduces linearly to the identity in \mathbb{F}_N . One should note that for reductions to 1 there is no difference between linear and cyclic reducibility; any word that can be reduced cyclically to 1 can also be reduced linearly to 1. If the reduced word w' is not the identity then the situation will be quite different for the two cases $N = 1$ and $N > 1$. For $N = 1$ any cyclic reduction can also be achieved in a linear way, but this is not the case for $N > 1$ any more. In particular, whereas the linear reduction of a word is always unique, this is not true any more for the cyclic reduction. For example, the word $u_1^{-1}u_2u_2^{-1}u_1^{-1}u_2^{-1}u_1$, which reduces linearly to $u_1^{-1}u_1^{-1}u_2^{-1}u_1$, has two different cyclic reductions, namely $u_1^{-1}u_2^{-1}$ and $u_2^{-1}u_1^{-1}$. Since one can think of the cyclic reduction as acting on the letters arranged on a circle, it is clear that any two cyclic reductions are related by a cyclic permutation. Thus the length of a cyclic reduction is well defined.

Recall that the string $\epsilon_1, \epsilon_2, \epsilon_3, \dots, \epsilon_n$ is dominating if

$$\begin{aligned} &\epsilon_1 \\ &\epsilon_1 + \epsilon_2 \\ &\epsilon_1 + \epsilon_2 + \epsilon_3 \\ &\vdots \\ &\epsilon_1 + \epsilon_2 + \epsilon_3 + \cdots + \epsilon_n \end{aligned}$$

are all strictly positive.

Let us translate this property to the word $u^{\epsilon_1}u^{\epsilon_2}u^{\epsilon_3} \cdots u^{\epsilon_n}$, with $u = u_1$. Starting with any word $l_1 \cdots l_n$ with $l_i \in \{u_1^{\pm 1}, \dots, u_N^{\pm 1}\}$ we let s_j be its j -th prefix, $s_j = l_1 \cdots l_j$. Then the dominating property of $\epsilon_1, \epsilon_2, \epsilon_3, \dots, \epsilon_n$ is equivalent to the fact that no prefix s_j ($1 \leq j \leq n$) of $u_1^{\epsilon_1}u_1^{\epsilon_2}u_1^{\epsilon_3} \cdots u_1^{\epsilon_n}$ is reducible to 1.

For example let $\epsilon_1 = \epsilon_2 = \epsilon_3 = 1$ and $\epsilon_4 = \epsilon_5 = -1$; then the only cyclic permutation of $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4, \epsilon_5$ that is dominating is the identity permutation. Equivalently, of the five cyclic permutations of $uuuu^{-1}u^{-1}$, only $uuuu^{-1}u^{-1}$ (i.e. apply the identity permutation) has no prefixes which are reducible to 1.

Now let us consider the same problem when the word contains more than one generator. We start with a word $w = l_1 \cdots l_n$ with $l_i \in \{u_1^{\pm 1}, \dots, u_N^{\pm 1}\}$. We say that w has *good reduction* if

- (i) no prefix of w is reducible to 1 and
- (ii) the linear reduction of w is cyclically reduced. (In the case of only one generator, this condition is always satisfied.)

We can then ask how many cyclic permutations of w have good reduction; and our main result is that this is the same as the number of letters in a cyclic reduction of w .

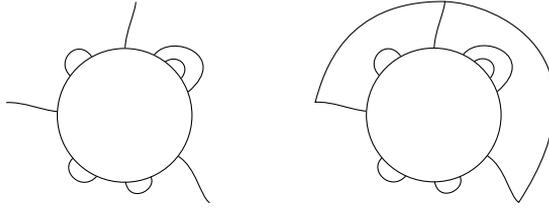


Fig. 1. On the left is π and on the right is $\tilde{\pi}$.

Theorem (The Non-Commutative Cycle Lemma). *Let w be a word and k the length of a cyclic reduction of w . Then w has exactly k cyclic permutations with good reduction.*

Formally, the Non-Commutative Cycle Lemma is also true for $k = 0$, because then no cyclic permutation has good reduction (as any cyclic permutation is reducible to 1). For $k \geq 1$, the statement is not so obvious because of the difference between linear and cyclic reducibility. We will in the following always assume that $k \geq 1$.

There is a way to say which letters of w remain after cyclic reduction – as we will see, a letter l_i of $w = l_1 \cdots l_n$ remains if the word $l_i l_{i+1} \cdots l_n l_1 \cdots l_{i-1}$ has good reduction. Moreover, we also show that one can canonically assign to each word of length n that cyclically reduces to a word of length k , a planar diagram, called a non-crossing circular half-pairing on $[n]$ with k through strings (see Fig. 5). Let us begin by recalling some definitions.

Let n be a positive integer and $[n] = \{1, 2, 3, \dots, n\}$. By a partition π of $[n]$ we mean a decomposition of $[n]$ into non-empty disjoint subsets $\pi = \{V_1, \dots, V_r\}$, i.e.

$$[n] = V_1 \cup \dots \cup V_r \quad \text{and} \quad V_i \cap V_j = \emptyset \quad \text{for } i \neq j$$

The subsets V_i are called the blocks of π , and we write $i \sim_\pi j$ if $i, j \in [n]$ are in the same block of π . We say that π has a crossing if we can find $i_1 < i_2 < i_3 < i_4 \in [n]$ such that

$$i_1 \sim_\pi i_3 \quad \text{and} \quad i_2 \sim_\pi i_4, \quad \text{but} \quad i_1 \not\sim_\pi i_2$$

We say that π is *non-crossing* if it has no crossings. A partition is called a *pairing* if all its blocks have exactly two elements; this can only happen when n is even. See [11] or [12] for a full discussion of non-crossing partitions.

We next wish to consider a special kind of non-crossing partition called a half-pairing.

Definition 1. Let π be a non-crossing partition in which no block has more than two elements and for which we have at least one block of size 1. From π create a new partition $\tilde{\pi}$ by joining into a single block all the blocks of π of size 1. If $\tilde{\pi}$ is non-crossing we say that π is a non-crossing *half-pairing*. The blocks of π of size 1 are called the *through strings* (see Fig. 1).

Note that we require a half-pairing to have at least one through string. This corresponds to $k \geq 1$ in our Cycle Lemma.

Let us relate this definition to our good reduction problem. Let $w = l_1 \cdots l_n$ with $l_i \in \{u_1^{\pm 1}, \dots, u_N^{\pm 1}\}$ be a word with n letters that cyclically reduces to a word of length k . We wish to assign to w a unique non-crossing half-pairing on $[n]$ with k through strings.

Definition 2. Let $w = l_1 \cdots l_n$ and π be a non-crossing half-pairing on $[n]$. We say that π is a *w-pairing* if

- (i) if (r, s) is a pair of π then $l_r = l_s^{-1}$ and
- (ii) if the singletons of π are $(i_1), (i_2), \dots, (i_k)$, then $l_{i_1} l_{i_2} \cdots l_{i_k}$ is a cyclic reduction of w .

Given w there may be more than one π which is a pairing of w . See Fig. 2.

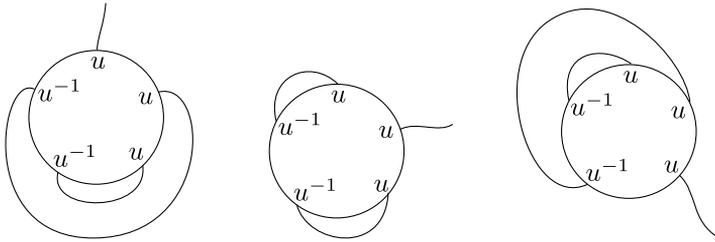


Fig. 2. The three possible w -pairings of $uuu^{-1}u^{-1}$.

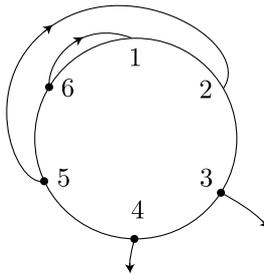


Fig. 3. The four outward oriented points are 3, 4, 5, and 6. 3 covers 4, 4 covers 5 and 3, and 5 covers 6.

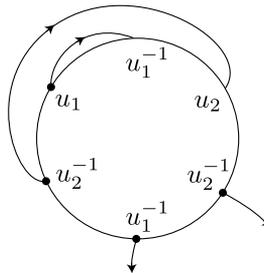


Fig. 4. Let $w = u_1^{-1}u_2u_2^{-1}u_1^{-1}u_2^{-1}u_1$ and $\pi = \{(1, 6), (2, 5), (3), (4)\}$. π is w -admissible. The second and third diagrams in Fig. 2 are not w -admissible.

In order to have a unique half-pairing associated with a word we impose a third condition which, in particular, will exclude the second and third diagrams in Fig. 2.

Definition 3. Let π be a non-crossing half-pairing of $[n]$. To each $i \in [n]$ we assign an orientation: out or in. Each singleton is assigned the *out* orientation. For each pair (r, s) of π exactly one of the cyclic intervals $[r, s]$ or $[s, r]$ contains a singleton (recall that we have at least one singleton). If $[r, s]$ does not contain a singleton, then we assign r the *out* orientation and s the *in* orientation.

Definition 4. Let π be a non-crossing half-pairing. We say that i covers the letter j if both have the out orientation and either $i + 1 = j$, or π pairs each letter of the cyclic interval $[i + 1, j - 1]$ with some other letter in the cyclic interval $[i + 1, j - 1]$. In particular this means that π has no singletons in $[i + 1, j - 1]$. (See Fig. 3.)

Definition 5. Let $w = l_1 \cdots l_n$ be a word and π a non-crossing half-pairing of $[n]$. We say that π is w -admissible if it is a w -pairing and we have for all i and j , $l_i \neq l_j^{-1}$ whenever i covers j . (See Fig. 4.)

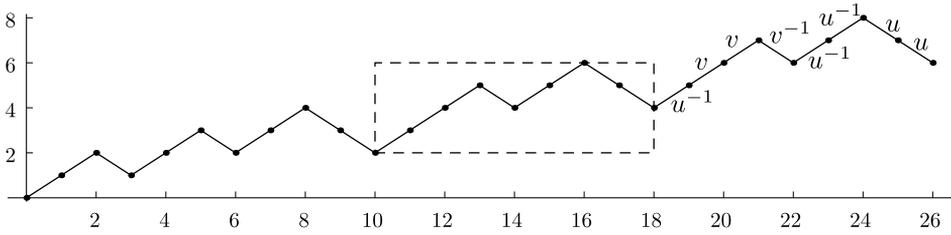


Fig. 6. Let $w = uuu^{-1}vvv^{-1}u^{-1}u^{-1} = x\tilde{w}y$ where $x = uu$, $\tilde{w} = u^{-1}v$, and $y = vv^{-1}u^{-1}u^{-1}$. The graph of t_i is shown, where t_i is the length of the linear reduction of the first i letters of w_∞ . The graph becomes shift-periodic at $i = 10$. The region enclosed in dotted lines shows one period.

reduces to the identity in \mathbb{F}_N . If $w = l_1 \cdots l_n$ is a word we let $w^{-1} = l_n^{-1} \cdots l_1^{-1}$; i.e. we reverse the string and take the inverse of each letter but do not do any reduction. Given a word w we let $|w|$ be the length of the linear reduction of w .

Remark 10. Let w be a word whose linear reduction is not cyclically reduced. Then either the first and last letter of w must cancel each other, or this cancellation must happen after removing a prefix or a suffix which is reducible to 1. By repeatedly cancelling letters at the ends of w and removing prefixes or suffixes which are reducible to 1, we are left with a word \tilde{w} which neither has a prefix or suffix which is reducible to 1 nor has cancellation of the first and last letters. For such a \tilde{w} the cyclic and linear reduction are the same. Thus for every word w there is a word \tilde{w} whose linear reduction is cyclically reduced and words x and y such that xy is reducible to 1 and such that we have as a concatenation of strings $w = x\tilde{w}y$. Depending on the order of cancellation different decompositions of a word may be found – we only require the existence of such a decomposition. See Fig. 6 for an example.

Proposition 11. Let w be a word of length n which reduces cyclically to a word of length $k > 0$. Let s be a prefix of w_∞ . If the number of letters in s exceeds $n(1 + n/k)$ then $|ws| = k + |s|$.

Proof. Write $w = x\tilde{w}y$ with x , y , and \tilde{w} as in Remark 10 above. For each positive integer m the length of the linear reduction of $(\tilde{w})^m$ is mk .

Now the linear reduction of x and the linear reduction of y are inverses of each other, so the last letter of the linear reduction of x is the inverse of the first letter of the linear reduction of y . Thus, if there is any cancellation between x and \tilde{w}^m there can be none between \tilde{w} and y , and vice versa for cancellation between \tilde{w}^m and y , i.e. if there is cancellation between \tilde{w} and y there can be none between x and \tilde{w} .

Let s be a prefix of w_∞ with $i > n(1 + n/k)$ letters. We shall show that $|ws| = k + |s|$. Let $m = \lceil i/n \rceil$. Since $i > n(1 + n/k)$, we have $m > n/k$.

First, suppose there is cancellation between x and \tilde{w} but none between \tilde{w} and y . Write s as $w^m s'$ with s' a prefix of w . Since $m > n/k$ the last letter in the linear reduction of $x\tilde{w}^m$ is the last letter of \tilde{w} . Thus $|s| = |x\tilde{w}^m y s'| = |x\tilde{w}^m| + |y s'|$ and likewise

$$\begin{aligned} |ws| &= |x\tilde{w}^{m+1} y s'| = |x\tilde{w}^{m+1}| + |y s'| \\ &= |x\tilde{w}^m| + |\tilde{w}| + |y s'| = k + |x\tilde{w}^m| + |y s'| \\ &= k + |s| \end{aligned}$$

Conversely, suppose that there is cancellation between \tilde{w} and y but none between x and \tilde{w} . Then, as $m > n/k$, the linear reduction of $\tilde{w}^m y$ begins with the same letter as does \tilde{w} . Hence $|s| = |x\tilde{w}^m y s'| = |x| + |\tilde{w}^m y s'|$ and

$$\begin{aligned}
 |ws| &= |x\tilde{w}^{m+1}ys'| = |x| + |\tilde{w}^{m+1}ys'| \\
 &= |x| + |\tilde{w}| + |\tilde{w}^m ys'| = k + |x| + |\tilde{w}^m ys'| \\
 &= k + |s| \quad \square
 \end{aligned}$$

Proof of the Non-Commutative Cycle Lemma. Let t_i be the length of the linear reduction of the first i letters of w_∞ . Choose $i_0 > n(1 + n/k)$. Then by Proposition 11, for any $i \geq i_0$, $t_{n+i} = k + t_i$. Choose i_1 to be the smallest $i_1 \geq i_0$ such that $t_{i_1} < t_j$ for all $j > i_1$, i.e. i_1 is the largest i such that $t_i = t_{i_0}$. Choose i_2 to be the largest i such that $t_{i_2} = 1 + t_{i_1}$. In general for $l \leq k$, choose i_l to be the largest i such that $t_{i_l} = l - 1 + t_{i_1}$. Since $t_{i_l-n} = t_{i_l} - k \leq t_{i_1}$, we must have $i_1 < i_2 < \dots < i_k \leq i_1 + n$.

For each l we choose the cyclic permutation of w that starts after the i_l -th letter of w_∞ . Since $t_{i_l+n} - t_{i_l} = k$ for such a word the linear and cyclic reductions are the same. Also since t never descends back to t_{i_l} , such a word will have no prefix which is reducible to 1. Thus it has good reduction.

If we choose a cyclic permutation at an i such that $t_i = t_{i-1} - 1$, the resulting word will be such that its linear reduction is not cyclically reduced. Indeed, let s be the prefix of w_∞ consisting of the first $i - 1$ letters, and let w_0 be the n letters following s . We must show that $|w_0| > k$. Since $t_i = t_{i-1} - 1$ there is cancellation between s and w_0 . Thus $t_{i-1} + |w_0| = |s| + |w_0| > |sw_0| = t_{i+n-1} = k + t_{i-1}$, hence $|w_0| > k$.

If we choose a cyclic permutation that starts at an i for which there is $j > i$ with $t_j = t_i$, the resulting word will also have a prefix which is reducible to 1. Thus there are only k cyclic permutations that produce good reduction. \square

Theorem 12. *Let w be a word of length n . Then there is a unique non-crossing half-pairing on $[n]$ which is w -admissible.*

Proof. Suppose $w = l_1 \dots l_n$ has good reduction then we construct the unique half-pairing which is w -admissible as follows. Starting with l_1 and moving to the right find the first $i < n$ such that $l_i = l_{i+1}^{-1}$. Pair these elements and return to l_1 and repeat the process, skipping over any letters already paired. Continue passing through w until no further pairings can be made. See Fig. 5. Put through strings on any unpaired letters. This produces a half-pairing which we denote π . As the pairs only involve adjacent letters or pairs that are adjacent after removing an adjacent pair no crossings will be produced. Moreover no pair (r, s) , $r < s$, will be produced with an unpaired letter in between. Thus the partition will be a non-crossing half-pairing.

To show that π is w -admissible we must show that there are no i and j such that i covers j and $l_i = l_j^{-1}$. Suppose i covers j , then according to the definition, both have the out orientation.

Let us break this into two cases. First case: i is a through string. Since i covers j , either $j = i + 1$ or π pairs every point of the cyclic interval $[i + 1, j - 1]$ with another point of $[i + 1, j - 1]$. In the first of these possibilities the algorithm would have paired i with j unless $i = n$, but this would imply that the linear and cyclic reduction of w are not the same. Thus we are left with the case that π pairs every point of the cyclic interval with another point in $[i + 1, j - 1]$.

Since w has good reduction, π starts with a through string – else w would have a prefix which is reducible to 1. Thus we must have $i < j$ since otherwise the cyclic interval $[i + 1, j - 1]$ would contain a through string. Hence each number in the interval $[i + 1, j - 1]$ is paired by π with another number in the interval $[i + 1, j - 1]$. Now our algorithm would have paired l_i with l_j , so we cannot have $l_i = l_j^{-1}$.

The second case is when i is the opening point of a pair (i, j') of π . We must have $i < j' < j$, for otherwise our algorithm would have paired i with j . However this contradicts our assumption that each number in the interval $[i + 1, j - 1]$ is paired by π with another in the interval. Thus π is w -admissible.

To see that π is unique notice that each time we add a pair it is a forced move. Indeed suppose i is the first i , starting from the left, such that $l_i = l_{i+1}^{-1}$. We cannot pair l_i with any earlier element because that would imply the earlier existence of an adjacent pair; we cannot pair l_i with any later letter as this would force i to cover $i + 1$. We then look for the next pair of elements either adjacent or

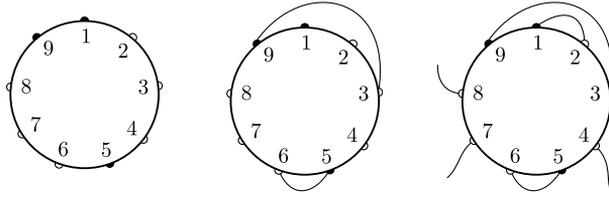


Fig. 7. At the left a dot diagram, in the centre the same diagram with a few strings added, and at the right the completed diagram.

adjacent after skipping over $\{i, i + 1\}$. By the same argument this pairing is also forced and continuing in this way we see that all pairs are forced and thus there is only one w -admissible half-pairing.

Now suppose that w does not have good reduction. By the Non-Commutative Cycle Lemma there are k cyclic permutations of w which have good reduction – one for each through string. Indeed, each cyclic permutation of w which has good reduction begins with a through string. Between each pair of through strings the method for pairing the elements is always the same: start at the first letter to the right of the through string and pair the first pair of adjacent letters that are inverses of each other and then return to the through string and repeat. Thus the method of pairing is entirely ‘local’ and is independent of at which through string we begin. \square

Lemma 13. *The number of non-crossing half-pairings on $[n]$ with k through strings is $\binom{n}{(n-k)/2}$.*

Proof. We use the method introduced in [7]. We place the points $1, 2, 3, \dots, n$ around the outside of a circle in clockwise order. On each point we shall place either a black dot or a white dot with a total of $(n - k)/2$ black dots and $(n + k)/2$ white dots. There are $\binom{n}{(n-k)/2}$ ways of doing this so we only have to show that each assignment of dots produces a unique non-crossing half-pairing and all half-pairings are produced in this way.

Now $(n - k)/2$ will be the number of pairs in the half pairing and each black dot will indicate which of the two points of the pair has the outward orientation. Starting at any black dot and moving clockwise search for the first available white dot not already paired with a black dot – except every time we pass over a black dot we skip a white dot to leave a white dot for the black dot to pair with. We proceed until all black dots are paired. Any remaining white dots become through strings.

Conversely starting with a non-crossing half-pairing on $[n]$ with k through strings, put a white dot on each through string and a white dot on the point of each pair with the inward orientation. Finally put a black dot on the point of each pair with the outward orientation. This gives the bijection between diagrams and dot patterns (see Fig. 7). \square

Theorem 14. *Let v be a cyclically reduced word of length k . The number of words of length n , whose standard cyclic reduction is v , is $(2N - 1)^{(n-k)/2} \times \binom{n}{(n-k)/2}$, in particular this number does not depend on v .*

Proof. Let $w = l_1 \dots l_n$ be a word of length n whose standard cyclic reduction is $v = v_1 \dots v_k$. By Theorem 12 there exists π , a unique non-crossing half-pairing π on $[n]$ with k through strings which is w -admissible. If the through strings of π are at i_1, \dots, i_k then $v_j = l_{i_j}$. We shall say that the j -th through string is coloured v_j . If (r, s) is a pair of π and r has the outward orientation then we shall say the pair (r, s) is coloured l_r . Thus each word whose standard cyclic reduction is v is associated with a unique non-crossing half-pairing coloured with the letters $\{u_1, u_1^{-1}, \dots, u_N, u_N^{-1}\}$ subject to the rule that no outward oriented point has the inverse colour of a point by which it is covered.

It remains to count how many of these coloured diagrams there are. By Lemma 13 there are $\binom{n}{(n-k)/2}$ diagrams with k through strings. The through strings are always coloured by the letters of v , so there is no choice here. However the outward oriented point of each pair can be coloured by any letter in $\{u_1, u_1^{-1}, \dots, u_N, u_N^{-1}\}$ except the inverse of the colour that covers it. Thus there are $2N - 1$ ways of choosing this colour. The colour of the inward oriented point of each pair is determined by

the colour of the corresponding outward oriented point of the pair. Thus once the diagram is selected there are $(2N - 1)^{\binom{n-k}{2}}$ ways of colouring it. \square

3. Concluding remarks

Suppose U_1, \dots, U_N are independent $m \times m$ Haar distributed random unitary matrices. Let u_1, \dots, u_N be the generators of the free group \mathbb{F}_N and $\mathbb{C}[\mathbb{F}_N]$ the group algebra of \mathbb{F}_N . Let $\phi: \mathbb{C}[\mathbb{F}_N] \rightarrow \mathbb{C}$ be the tracial linear functional defined on words in \mathbb{F}_N by $\phi(e) = 1$, $\phi(w) = 0$ for words $w \neq e$ and then extended to all of $\mathbb{C}[\mathbb{F}_N]$ by linearity. Suppose Y is a linear combination of words in $\{U_1, U_1^{-1}, \dots, U_N, U_N^{-1}\}$ and y is the corresponding linear combination of words in $\{u_1, u_1^{-1}, \dots, u_N, u_N^{-1}\}$. Voiculescu [14] showed that

$$\lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{1}{m} \text{Tr}(Y) \right] = \phi(y),$$

thus establishing the asymptotic $*$ -freeness of the U_1, \dots, U_N . In particular this implies the asymptotic freeness of the self-adjoint operators X_1, \dots, X_N , where $X_i = U_i + U_i^{-1}$.

In recent years the fluctuation of random matrices has been the object of much study (see [2,7–10, 13]). Let $X = X_1 + \dots + X_N$ and for integers p and q consider the asymptotic fluctuation moments:

$$\alpha_{p,q} = \lim_n \mathbb{E} \left[(\text{Tr}(X^p) - \mathbb{E}[\text{Tr}(X^p)]) \cdot (\text{Tr}(X^q) - \mathbb{E}[\text{Tr}(X^q)]) \right]$$

One way to understand these moments is via the theory of orthogonal polynomials. In this situation it means finding a sequence of polynomials $\{R_k\}_k$ such that for $k \neq l$ we have

$$\lim_n \mathbb{E} \left[(\text{Tr}(R_k(X)) - \mathbb{E}[\text{Tr}(R_k(X))]) \cdot (\text{Tr}(R_l(X)) - \mathbb{E}[\text{Tr}(R_l(X))]) \right] = 0$$

Such a sequence of polynomials is said to diagonalize the fluctuations. This has been done for a variety of random matrix ensembles (see [5] and [7] and the references there).

Corollary 8 suggests that there ought to be polynomials $\{P_n\}_n$ such that $\widehat{P}_n(x) = Q_n$. Indeed, using the Non-Commutative Cycle Lemma we have shown that the polynomials indicated by Corollary 8 do diagonalize the fluctuations of the operator X above. The polynomials can also be obtained by modifying the Chebyshev polynomials of the first kind: let $R_0(x) = 2$, $R_1(x) = 1$, and $R_{k+1}(x) = xR_k(x) - (2N - 1)R_{k-1}(x)$. Moreover for n odd $P_n = R_n$ and for n even $P_n(x) = R_n(x) + 2$. The proofs of these results will be presented in a subsequent paper.

References

- [1] J. Bertrand, Calcul des probabilités. Solution directe du problème du scrutin, C. R. Math. Acad. Sci. Paris 105 (1887) 369.
- [2] B. Collins, J.A. Mingo, P. Śniady, R. Speicher, Second order freeness and fluctuations of random matrices: III. Higher order freeness and free cumulants, Doc. Math. 12 (2007) 1–70.
- [3] A. Dvoretzky, Th. Motzkin, A problem of arrangements, Duke Math. J. 14 (1947) 305–313.
- [4] N. Dershowitz, S. Zaks, The Cycle Lemma and some applications, European J. Combin. 11 (1990) 35–40.
- [5] K. Johansson, On fluctuations of eigenvalues of random Hermitian matrices, Duke Math. J. 91 (1998) 151–204.
- [6] H. Kesten, Symmetric random walks on groups, Trans. Amer. Math. Soc. 92 (1959) 336–354.
- [7] T. Kusalik, J.A. Mingo, R. Speicher, Orthogonal polynomials and fluctuations of random matrices, J. Reine Angew. Math. 604 (2007) 1–46.
- [8] J. Mingo, A. Nica, Annular noncrossing permutations and partitions, and second-order asymptotics for random matrices, Int. Math. Res. Not. IMRN (28) (2004) 1413–1460.
- [9] J. Mingo, P. Śniady, R. Speicher, Second order freeness and fluctuations of random matrices: II. Unitary random matrices, Adv. Math. 209 (2007) 212–240.
- [10] J. Mingo, R. Speicher, Second order freeness and fluctuations of random matrices: I. Gaussian and Wishart matrices and cyclic Fock spaces, J. Funct. Anal. 235 (2006) 226–270.
- [11] A. Nica, R. Speicher, Lectures on the Combinatorics of Free Probability, London Math. Soc. Lecture Note Ser., vol. 335, Cambridge University Press, 2006.
- [12] R. Simion, Noncrossing partitions, Discrete Math. 217 (2000) 367–409.
- [13] R. Rao, J. Mingo, A. Edelman, R. Speicher, Statistical eigen-inference from large Wishart matrices, Ann. Statist. 36 (2008) 2850–2885.
- [14] D. Voiculescu, K. Dykema, A. Nica, Free Random Variables, CRM Monogr. Ser., vol. 1, Amer. Math. Soc., 1992.