Extending from bijections between marked occurrences of patterns to all occurrences of patterns

Jeffrey Remmel and Mark Tiefenbruck

Department of Mathematics, University of California, San Diego, USA

Abstract. We consider two recent open problems stating that certain statistics on various sets of combinatorial objects are equidistributed. The first, posed by Anders Claesson and Svante Linusson, relates nestings in matchings on $\{1, 2, \ldots, 2n\}$ to occurrences of a certain pattern in permutations in S_n. The second, posed by Miles Jones and Jeffrey Remmel, relates occurrences of a large class of consecutive permutation patterns to occurrences of the same pattern in the cycles of permutations. We develop a general method that solves both of these problems and many more. We further employ the Garsia-Milne involution principle to obtain purely bijective proofs of these results.

Résumé. Nous considérons deux dernières problèmes ouverts indiquant que certaines statistiques sur les divers ensembles d'objets combinatoires sont équiréparties. La première, posée par Anders Claesson et Svante Linusson, concerne les imbrications dans des filtrages sur $\{1, 2, \ldots, 2n\}$ pour les occurrences d'un certain modèle de permutations dans S_n. La seconde, posée par Miles Jones et Jeffrey Remmel, concerne les occurrences d'une large classe de schémas de permutation consécutive aux événements du même modèle dans les cycles de permutations. Nous développons une méthode générale qui résout ces deux problèmes et beaucoup plus. Nous avons également utiliser le principe d'involution Garsia-Milne pour obtenir des preuves purement bijective de ces résultats.

Keywords: bijection, permutation statistics, generating function, partially marked pattern family

1 Introduction

In this paper, we present a general method for proving that certain pairs of statistics are equidistributed. In short, we embed the objects under consideration into much larger sets. Given a bijection between these larger sets with certain properties, our method can be used to prove the original statement. Further, our method may be augmented with the Garsia-Milne involution principle [3] to obtain bijective proofs of these results.

As examples of the method, we consider two recent open problems, one by Claesson and Linusson [2] and one by Jones and Remmel [4]. In the remainder of this section, we introduce these two problems. In Section 2, we present the general method by which these and many other problems can be solved. Finally, in Section 3, we provide the necessary bijections for these two problems.

1.1 Claesson and Linusson's problem

Define the set $[n] = \{1, 2, \ldots, n\}$. Following Claesson and Linusson [2], we make the following definitions. A *matching* on $[2n]$ is a partition of that set into blocks of size 2. An example of a matching is

$$M = \{(1,3),(2,7),(4,6),(5,8)\}. \tag{1}$$

In the diagram below there is an *arc* connecting i with j precisely when $(i,j) \in M$.

A *nesting* of M is a pair of arcs (i,ℓ) and (j,k) with $i < j < k < \ell$:

We call such a nesting a *left-nesting* if $j = i + 1$. Similarly, we call it a *right-nesting* if $\ell = k + 1$. A nesting may be a left-nesting, a right-nesting, both, or neither. M has one nesting, formed by the arcs $(2,7)$ and $(4,6)$. It is a right-nesting.

In a permutation $\pi = \pi_1 \ldots \pi_n$, an occurrence of the pattern

is a subsequence $\pi_i \pi_{i+1} \pi_j$ such that $\pi_j + 1 = \pi_i < \pi_{i+1}$. As an example, the permutation $\pi = 351426$ contains one such occurrence, 352. It can be seen graphically in the following permutation diagram:

Bousquet-Mélou et al. [1] gave bijections between matchings on $[2n]$ with no left- or right-nestings and three other classes of combinatorial objects, thus proving that they are equinumerous. The other classes were unlabeled $(\mathbf{2}+\mathbf{2})$-free posets (or interval orders) on n nodes; permutations in S_n avoiding the pattern ; and ascent sequences of length n. Claesson and Linusson conjectured that the distribution of right-nestings in matchings on $[2n]$ with no left-nestings is equal to the distribution of occurrences of in S_n.

1.2 Jones and Remmel's problem

Let $\sigma = \sigma_1 \sigma_2 \ldots \sigma_n \in S_n$, let $w = (w_1 w_2 \ldots w_n)$ be a cycle in a permutation, and let $\pi \in S_m$, where $m \le n$. A *match* of the pattern π in σ is a sequence of consecutive elements $\sigma_i \sigma_{i+1} \ldots \sigma_{i+m-1}$ that have the same relative order as the entries of π. A *cycle-match* of the pattern π in w is a sequence of consecutive elements of w, where w_1 follows w_n, that have the same relative order as the entries in π. For example, in the permutation 35124, 351 is a 231-match, and in the cycle (34769), 693 is a 231-cycle-match.

Jones and Remmel [4] showed that if $\pi \in S_m$ begins or ends with 1 or m, then the number of π-cycle-matches in the cycles of the permutations in S_n has the same distribution as the number of π-matches in the permutations in S_n. They conjectured this was true for any π that cannot cover a cycle with overlapping π-cycle-matches. For example, in the cycle $w = (31425)$, 3142 and 4253 are 3142-cycle-matches that cover w, whereas no cycle can be covered by 2143-cycle-matches.

2   Main Theorem

Direct arguments to prove these two conjectures remain elusive. Instead, we consider the set of objects where each occurrence is either "marked" or "not marked". For example, in extending the set of matchings on $[2n]$, a matching with k right-nestings would appear 2^k times in the new set, once with each subset of its right-nestings considered marked. We wish to find a bijection θ_n that takes a matching M on $[2n]$ with no left-nestings and sends it to a permutation $\sigma \in S_n$ with the same number of occurrences of ⊞ as M has right-nestings. Instead, our method is to produce a bijection Γ_n between these larger sets that preserves the number of *marked* occurrences. In this section, we give a very general description of this setup and state the desired result.

The main purpose of this paper is to show that there is a general mechanism to construct our desired bijections θ_n from the bijections Γ_n. We consider the following setup. Let $\mathbb{N}$ denote the set of natural numbers. For any alphabet A, we let A^* denote the set of all words over A and let A^n denote the set of all words of length n over A. We let ϵ denote the empty word. A *pattern family* is a set of the form $\mathcal{F} = \bigcup_{n \geq 0} \langle f, g, \mathcal{F}_n, A_n, \mathcal{P}_n \rangle$ where f and g are functions mapping $\mathbb{N}$ into $\mathbb{N}$ and for each $n \geq 0$,

1. A_n is a finite alphabet,

2. $\mathcal{F}_n$ is a subset of $A_n^{f(n)}$,

3. $\mathcal{P}_n$ is a totally ordered set of *patterns*, $P_1, P_2, \ldots, P_{g(n)}$, where for each i, P_i is a set of pairs $\langle a_1 a_2 \ldots a_k, T_{a_1 a_2 \ldots a_k} \rangle$ where $1 \leq a_1 < \cdots < a_k \leq f(n)$ and $T_{a_1 a_2 \ldots a_k} \subseteq A_n^k$. Here we do not require that the sequences $a_1 a_2 \ldots a_k$ that appear as the first elements of pairs in P_i are all of the same length, but we do insist that all such sequences are distinct.

An *occurrence* of the pattern $P_i \in \mathcal{P}_n$ in a sequence $s = s_1 \ldots s_{f(n)} \in \mathcal{F}_n$ is a subsequence of indices $a_1 a_2 \ldots a_k$ such that there exists a pair $\langle a_1 a_2 \ldots a_k, T_{a_1 a_2 \ldots a_k} \rangle \in P_i$ and $s_{a_1} s_{a_2} \ldots s_{a_k} \in T_{a_1 a_2 \ldots a_k}$. We let $P_i(s)$ denote the number of occurrences of the pattern P_i in the sequence s. We shall be interested in the generating function

$$R_{\mathcal{F}}(t, x_1, x_2, \ldots) = \sum_{n \geq 0} t^n \sum_{s \in \mathcal{F}_n} \prod_{i=1}^{g(n)} x_i^{P_i(s)} \tag{2}$$

as well as its specialization

$$R_{\mathcal{F}}(t, x) = \sum_{n \geq 0} t^n \sum_{s \in \mathcal{F}_n} x^{\sum_{i=1}^{g(n)} P_i(s)} \tag{3}$$

We note that we allow $g(n) = 0$, in which case we will assume that $\mathcal{P}_n = \emptyset$ and interpret $\prod_{i=1}^{g(n)} x_i^{P_i(s)}$ and $x^{\sum_{i=1}^{g(n)} P_i(s)}$ to be equal to 1.

Given a pattern family $\mathcal{F} = \bigcup_{n \geq 0} \langle f, g, \mathcal{F}_n, A_n, \mathcal{P}_n \rangle$, we can form the *partially marked pattern family* $\mathcal{PMF} = \bigcup_{n \geq 0} \langle f, g, \mathcal{PMF}_n, A_n, \mathcal{P}_n \rangle$ from $\mathcal{F}$ where if $\mathcal{P}_n = \{P_1, \ldots, P_{g(n)}\}$, then $\mathcal{PMF}_n$ consists of all $(g(n) + 1)$-tuples $\langle s, H_1, \ldots, H_{g(n)} \rangle$ such that $s \in \mathcal{F}_n$ and for $i = 1, \ldots, g(n)$, H_i is a possibly empty set of occurrences of the pattern P_i in s. Thus we can think of the $(g(n) + 1)$-tuple $\langle s, H_1, \ldots, H_{g(n)} \rangle$ as an element $s \in \mathcal{F}_n$ where some of the occurrences of P_i in s are "marked" for $i = 1, \ldots, g(n)$. We define the weight of $\langle s, H_1, \ldots, H_{g(n)} \rangle$ to be

$$w_{\mathcal{PMF}}(s, H_1, \ldots, H_{g(n)}) = \prod_{i=1}^{g(n)} y_i^{|H_i|} \tag{4}$$

where again we make the convention that if $g(n) = 0$, then we set $w_{\mathcal{PMF}}(s) = 1$. In such a situation, we shall consider the generating function

$$MR_{\mathcal{F}}(t, y_1, y_2, \ldots) = \sum_{n \geq 0} t^n \sum_{(s, H_1, \ldots, H_{g(n)}) \in \mathcal{PMF}_n} w_{\mathcal{PMF}}(s, H_1, \ldots, H_{g(n)}) \tag{5}$$

as well as its specialization

$$MR_{\mathcal{PMF}}(t, y) = MR_{\mathcal{F}}(t, y, y, \ldots) \tag{6}$$

The key result of this paper is the following theorem.

Theorem 1 *Suppose that $\mathcal{F}$ is a pattern family and $\mathcal{PMF}$ is the partially marked pattern family constructed from $\mathcal{F}$. Then*

$$MR_{\mathcal{F}}(t, x_1 - 1, x_2 - 1, \ldots) = R_{\mathcal{F}}(t, x_1, x_2, \ldots), \tag{7}$$

so that

$$MR_{\mathcal{F}}(t, x - 1) = R_{\mathcal{F}}(t, x). \tag{8}$$

We prove Theorem 1 via a simple involution. Now Theorem 1 has the following obvious corollary.

Corollary 1 *Suppose that $\mathcal{F} = \bigcup_{n \geq 0} \langle f, g, \mathcal{F}_n, A_n, \mathcal{P}_n \rangle$ and $\mathcal{G} = \bigcup_{n \geq 0} \langle h, g, \mathcal{G}_n, B_n, \mathcal{Q}_n \rangle$ are pattern families. (Here we are not insisting that $f = h$ which means that for any given n, the elements of $\mathcal{F}_n$ and $\mathcal{G}_n$ can have different lengths, but we are insisting that the number of patterns in $\mathcal{P}_n$ and $\mathcal{Q}_n$ are the same.) Let $\mathcal{PMF} = \bigcup_{n \geq 0} \langle f, g, \mathcal{PMF}_n, A_n, \mathcal{P}_n \rangle$ and $\mathcal{PMG} = \bigcup_{n \geq 0} \langle h, g, \mathcal{PMG}_n, B_n, \mathcal{Q}_n \rangle$ be the partially marked pattern families constructed from $\mathcal{F}$ and $\mathcal{G}$, respectively. Then*

$$MR_{\mathcal{F}}(t, y_1, y_2, \ldots) = MR_{\mathcal{G}}(t, y_1, y_2, \ldots) \tag{9}$$

implies

$$R_{\mathcal{F}}(t, x_1, x_2, \ldots) = R_{\mathcal{G}}(t, x_1, x_2, \ldots). \tag{10}$$

We can give a completely bijective proof of Corollary 1 by combining our proof of Theorem 1 with the involution principle of Garsia and Milne [3]. That is, for all $n \geq 0$, if there is a bijection $\Gamma_n : \mathcal{PMF}_n \to \mathcal{PMG}_n$ such that for all $(s, H_1, \ldots, H_{g(n)}) \in \mathcal{PMF}_n$,

$$w_{\mathcal{PMF}}(s, H_1, \ldots, H_{g(n)}) = w_{\mathcal{PMG}}(\Gamma_n((s, H_1, \ldots, H_{g(n)}))), \tag{11}$$

then we can construct a bijection $\theta_n : \mathcal{F}_n \to \mathcal{G}_n$ from Γ_n such that for all $s \in \mathcal{F}_n$,

$$\prod_{i=1}^{g(n)} x_i^{P_i(s)} = \prod_{i=1}^{g(n)} x_i^{Q_i(\theta_n(s))} \tag{12}$$

where $\mathcal{P}_n = \{P_1, \ldots, P_{g(n)}\}$ and $\mathcal{Q}_n = \{Q_1, \ldots, Q_{g(n)}\}$.

3   Bijections between marked occurrences

With Corollary 1 in hand, all we need is to define the bijections Γ_n that preserve the number of marked occurrences. In this section, we provide those bijections for the two problems stated in Section 1.

3.1   *Claesson and Linusson's problem*

In order to use Corollary 1, we first must express our sets as pattern families. While this is straightforward for occurrences of ▦ in S_n, matchings on $[2n]$ do not easily fit into this framework. We will first need to define an encoding of matchings with no left-nestings as words, then define a pattern that corresponds to right-nestings. There are a few obvious encodings, but Claesson and Linusson [2] give a bijection between these matchings and inversion tables that will better suit our needs.

Let $\mathcal{M}_n$ be the set of matchings on $[2n]$, and let $M \in \mathcal{M}_n$. If $i < j$ and $\alpha = (i, j)$ is an arc in M, then we call i the *opener* of α and j the *closer* of α. Order the arcs of M by closer, so for example, the arc with closer $2n$ is the nth arc. Define $\mathcal{I}_n$ as the Cartesian product

$$\mathcal{I}_n = [0,0] \times [0,1] \times \cdots \times [0, n-1], \tag{13}$$

where $[i, j] = \{i, i+1, \ldots, j\}$. An element of $\mathcal{I}_n$ is called an *inversion table*. Then the function $f : \mathcal{M}_n \to \mathcal{I}_n$, given by $f(M) = (a_1, a_2, \ldots, a_n)$, where a_i is the number of closers less than the opener of the ith arc of M, is a bijection. Tracing the definition through, we find that a right nesting corresponds to an index $k < n$ such that $a_k > a_{k+1}$ and $a_i \neq k$ for all $i > k+1$. We call k a *proscriptive descent*.

Clearly, $\mathcal{I}_n$ can be written as a set of words on the alphabet $[0, n-1]$. Moreover, we can construct a pattern family corresponding to proscriptive descents. We will now use this setup to construct the bijection Γ_n. Our construction will be recursive, so define Γ_0 to map the empty inversion table to the empty permutation, and suppose that Γ_{n-1} has already been defined. Let $I = (a_1, a_2, \ldots, a_{n-1}) \in \mathcal{I}_{n-1}$, and let H be a set of proscriptive descents in I. Let $(\sigma, J) = \Gamma_{n-1}(I, H)$ be the corresponding permutation in S_{n-1} and set of occurrences of ▦.

We consider the ways that I can be extended to an element of $\mathcal{I}_n$ with the same set H: a_n may be chosen freely from $[0, n-1] \setminus H$. Also, σ can be extended to an element of S_n with essentially the same set J by inserting a new smallest element ("essentially" because every element of σ is increased by 1 and some move to the right). For example, inserting into the second position of 35124 yields 416235. If $j > i+1$ and $\sigma_j + 1 = \sigma_i < \sigma_{i+1}$ is an occurrence of ▦ in J, then we may not insert a new smallest element at position $i+1$. Thus, there are $n - |J|$ places we may insert a new smallest element.

For $k = 1, 2, \ldots, n - |H|$, let $I^{(k)}$ be the extension of I such that a_n is chosen to be the kth smallest valid choice, let $\sigma^{(k)}$ be the extension of σ where a new smallest element is placed in the kth valid position from the right, and let $J^{(k)}$ be the corresponding set of occurrences of ▦. Then, define $\Gamma_n(I^{(k)}, H) = (\sigma^{(k)}, J^{(k)})$. Further, if $n-1$ is a proscriptive descent in $I^{(k)}$, then the new smallest element will create an occurrence of ▦ in $\sigma^{(k)}$ with $\sigma_i = 2$ and $\sigma_j = 1$. If $J'^{(k)}$ is the set $J^{(k)}$ with this occurrence included, then define $\Gamma_n(I^{(k)}, H \cup \{n-1\}) = (\sigma^{(k)}, J'^{(k)})$.

3.2   *Jones and Remmel's problem*

Fix π so that no cycle can be covered by overlapping π-cycle-matches. It is clear that π-matches and π-cycle-matches in permutations in S_n can be converted to the framework of pattern families. For $\sigma \in S_n$

and H a set of π-cycle-matches in σ, we need to define $\Gamma_n(\sigma, H) = (\tau, J)$ such that J is a set of π-matches in τ.

First, order the cycles of σ in decreasing order by smallest element. Write each cycle so that the smallest element is written first. If any π-cycle-match in H wraps around the end of its cycle, then rewrite the cycle so that it begins with that π-cycle-match. Repeat until no element of H wraps around the end of its cycle. This is always possible because of the condition on π. Finally, remove the parentheses from the cycles of σ to obtain τ, and set $J = H$. For example, if $\pi = 2143$, $\sigma = (1542)(36)$, and $H = \{2154\}$, then we first write $\sigma = (36)(1542)$, then $\sigma = (36)(2154)$, then $\tau = 362154$. On the other hand, if $H = \emptyset$, then $\tau = 361542$.

In this case, our bijection tells us even more: there is a bijection θ_n such that $\theta_n(\sigma)$ has the exact same π-matches as σ's π-cycle-matches. We may also extend this result in a different way: If P is a set of permutations such that no combination of them can cover a cycle with cycle-matches, then there is a bijection θ_n such that $\theta_n(\sigma)$ has the exact same matches as σ's cycle-matches among P. The following are two interesting classes of such sets.

If $\sigma \in S_n$, define $\sigma.^* = \{\tau : \tau \in \cup_{m \geq n} S_m, \tau_i = \sigma_i \text{ for } 1 \leq i \leq n\}$ and $\sigma.^+ = \{\tau : \tau \in \cup_{m > n} S_m, \tau_i = \sigma_i \text{ for } 1 \leq i \leq n\}$. Then, for any k,

$$k(k-1)...1.^+ \cup 12.^* \cup 213.^* \cup \cdots \cup (k-1)(k-2)...1k.^* \tag{14}$$

and

$$\bigcup_{\sigma \in S_k} \sigma.^+ \tag{15}$$

satisfy the necessary conditions.

4 Conclusion

In addition to the examples in this paper, we have found interesting new results involving rook placements and permutation statistics. It is clear that the method can be applied to many more types of combinatorial objects as well. We believe this is a powerful method of finding bijections that can be used to solve a wide variety of problems.

References

[1] M. Bousquet-Mélou, A. Claesson, M. Dukes and S. Kitaev, $(2 + 2)$-free posets, ascent sequences and pattern avoiding permutations, *Journal of Combinatorial Theory Series A* 117 (2010) 884–909.

[2] A. Claesson and S. Linusson, n! matchings, n! posets, *Proceedings of the American Mathematical Society* 139 (2011) 435–449

[3] A. Garsia and S. Milne, Method for constructing bijections for classical partition identities, *Proc. Nat. Acad. Sci. USA* 78 (1981), 2026–2028.

[4] M. Jones and J. Remmel, Pattern Matching in the Cycle Structures of Permutations