

Perfect sampling algorithm for Schur processes

D. Betea,^{*} C. Boutillier,^{*} J. Bouttier,[†] G. Chapuy,[‡] S. Corteel,[‡] M. Vuletić[§]

Abstract

We describe random generation algorithms for a large class of random combinatorial objects called *Schur processes*, which are sequences of random (integer) partitions subject to certain interlacing conditions. This class contains several fundamental combinatorial objects as special cases, such as plane partitions, tilings of Aztec diamonds, pyramid partitions and more generally steep domino tilings of the plane. Our algorithm, which is of polynomial complexity, is both *exact* (i.e. the output follows exactly the target probability law, which is either Boltzmann or uniform in our case), and *entropy optimal* (i.e. it reads a minimal number of random bits as an input). It can be viewed as a (far reaching) common generalization of the RSK algorithm for plane partitions and of the *domino shuffling algorithm* for domino tilings of the Aztec diamond. At a technical level, it relies on unified bijective proofs of the different types of Cauchy identities for Schur functions, and on an adaptation of Fomin’s growth diagram description of the RSK algorithm to that setting. Simulations performed with this algorithm suggest previously unobserved phenomena in the limit shapes for some tiling models.

1 Introduction

Tilings of the plane by pieces of prescribed shapes (such as dominos or rhombi) are fundamental combinatorial objects that have received much attention in discrete mathematics and computer science. In particular, several tiling problems have been studied as models of two dimensional statistical physics, mainly because their remarkable combinatorial structure

makes these models physically interesting, algorithmically manageable, and mathematically tractable.

The most celebrated example, introduced in [5], is given by domino tilings of the *Aztec diamond*, see Figure 1. The first remarkable property of this model is enumerative: the number of domino tilings of the Aztec diamond of size n is $2^{\frac{n(n+1)}{2}}$. This property was proved in [5] in several ways, but one of them is of particular importance for the present paper: this result can be proved using a bijective procedure, called the *domino shuffling algorithm*, that generates a tiling of the Aztec diamond of size n taking exactly $\frac{n(n+1)}{2}$ bits as an input. This algorithm is not only elegant but also very useful, since it enabled the efficient sampling of large random tilings, and led to the empirical discovery of the *arctic circle phenomenon* later proved in [7] (see Figure 7).

The purpose of this paper is to extend this picture to much more general models than the Aztec diamond. Indeed in recent years, many other models of tilings or related objects have been introduced and studied, mainly under the enumerative or “limit shape” viewpoint. In particular, it was recently observed that tilings of the Aztec diamond are a part of a much larger family of models of domino tilings of the plane, called *steep tilings* [3], see e.g. Figure 5. Even more, steep tilings can themselves be represented as sequences of integer partitions known as *Schur processes*, which puts them under the same roof as other well known objects such as *plane partitions* (Figure 7). The latter have been much studied as well, and a very elegant and efficient




Figure 1: Two domino tilings of the Aztec diamond of size 4.

^{*}Laboratoire de Probabilités et Modèles Aléatoires, UMR 7599, Université Pierre et Marie Curie, 4 place Jussieu, F-75005 Paris, dan.betea@gmail.com|cedric.boutillier@upmc.fr

[†]Institut de Physique Théorique, CEA, IPhT, 91191 Gif-sur-Yvette, France, CNRS URA 2306 and Département de Mathématiques et Applications, École normale supérieure, 45 rue d’Ulm, F-75231 Paris Cedex 05, jeremie.bouttier@cea.fr

[‡]LIAFA, CNRS et Université Paris Diderot, Case 7014, F-75205 Paris Cedex 13, (guillaume.chapuy|corteel)@liafa.univ-paris-diderot.fr

[§]Department of Mathematics, University of Massachusetts Boston, Boston, MA 02125, USA, mirjana.vuletic@umb.edu

strategy of enumeration and random generation for these objects follows from the Robinson-Schensted Knuth (RSK) algorithm, see [16, Chap 7].

Our main result is a random generation algorithm for general Schur processes. Remarkably, the algorithm is a common generalization of both the RSK based strategy for plane partitions and of the domino shuffling algorithm. It enables one to efficiently sample large random configurations, which at the empirical level unveils some new properties of their limit shapes (see Figure 8).

At a technical level, our main tools are combinatorial constructions dealing with integer partitions subject to certain interlacing conditions, that can be viewed as bijective proofs of the different types of Cauchy identities for Schur functions. The obtained algorithm takes as an input a product sequence of geometric random variables and random bits that can be represented in a graphical way in terms of *growth diagrams* similar to Fomin’s description of the RSK algorithm (see Fomin’s appendix in [16]). In this setting, it is easy to see that the algorithm is entropy optimal (Proposition 2). The (polynomial) complexity is also studied (Proposition 1).

We conclude this discussion with references to previous works. Borodin [1] gives a random sampling algorithm for Schur processes, but it is not entropy optimal (although this question is not discussed in [1], this is easy to see by comparing with the present paper). Propp and Wilson’s *coupling from the past* method [15] is often used for random sampling of domino tilings. In a different direction, the reader interested in exhaustive sampling of tilings (outside the realm of Schur processes) could consult [4] and references therein.

The paper is organized as follows: in Section 2 we describe the necessary prerequisites on partitions, Schur functions and the vertex operator formalism needed in the rest of the paper. In Section 3 we define the Schur process and give its relation to tilings. In Section 4 we give the main sampling algorithm for the Schur process after initially describing the two main bijections we use in said algorithm. We also give some samples of large tilings obtained using this algorithm. We conclude in Section 5.

2 Partitions, Schur functions and vertex operators

A partition λ is either the empty partition (0) or a sequence of strictly positive numbers listed in decreasing order: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0$. We

call each λ_i a *part* and $\ell(\lambda) := k$ the *length* (number of non-zero parts) of λ . We reserve letters λ, μ, \dots for partitions. We call $|\lambda| := \sum_{i=1}^{\ell(\lambda)} \lambda_i$ the *weight* of the partition. For any λ we have a *conjugate partition* λ' whose parts are defined as $\lambda'_i := |\{j : \lambda_j \geq i\}|$. We finally define the notion of interlacing partitions. Let λ and μ be two partitions with $\mu \subseteq \lambda$ (that is, $\mu_i \leq \lambda_i$ for all i). They are said to be *interlacing* and we write $\lambda \succ \mu$ if and only if

$$\lambda_1 \geq \mu_1 \geq \lambda_2 \geq \mu_2 \geq \lambda_3 \geq \dots$$

In the language of [11], the interlacing property is equivalent to saying that the skew diagram $\lambda - \mu$ forms a *horizontal strip*. There is a *dual interlacing notion*. We write $\mu \prec' \lambda$ if the skew diagram $\lambda - \mu$ forms a *vertical strip* ($0 \leq \lambda_i - \mu_i \leq 1$ for all i). Hereinafter, we will the denote the partition (m, \dots, m) (n parts) by m^n .

We will represent partitions graphically by either Young or (when convenient) Maya diagrams. A *Maya diagram* is an encoding of the Young diagram as a boolean function $\mathbb{Z} + \frac{1}{2} \rightarrow \{\circ, \bullet\}$. See Figure 2.




Figure 2: Young diagram of partition $(2, 2, 2, 1, 1)$ (left, English notation) and the Maya diagram associated to it (right, Russian notation for the partition; for French notation, see below). Far enough to the left there are only holes and far enough to the right, only particles.

Partitions have many uses in combinatorics, statistical mechanics, and other fields. For our purposes, they appear in the study of the algebra of symmetric functions. Partitions of length less than or equal to n index any linear basis of $\mathbb{C}[x_1, \dots, x_n]^{S_n}$. We will be interested in the basis of Schur polynomials. They can be defined by, for example, the Jacobi-Trudi formula (see [11]):

$$s_\lambda(x_1 \dots x_n) = \det_{1 \leq i, j \leq n} h_{\lambda_i - i + j}(x_1 \dots x_n)$$

where the h 's are the complete symmetric polynomials and $s_0(\cdot) = 1$, $s_\lambda(\cdot) = \delta_{\lambda, 0}$. Related are the skew Schur functions, which for two partitions λ, μ are zero unless $\mu \subseteq \lambda$, in which case they are defined by:

$$s_{\lambda/\mu}(x_1 \dots x_n) = \det_{1 \leq i, j \leq n} h_{\lambda_i - \mu_j - i + j}(x_1 \dots x_n)$$

Notice $s_{\lambda/0} = s_\lambda$. In one variable, $s_{\lambda/\mu}(x_1) = x_1^{|\lambda - \mu|} \delta_{\lambda \succ \mu}$. Let us denote the alphabet $(x_1 \dots x_n)$

by X (and similarly for Y). For our purposes, the important identities satisfied by Schur functions are the branching rule

$$(1) \quad s_{\lambda/\mu}(X, Y) = \sum_{\nu} s_{\lambda/\nu}(X) s_{\nu/\mu}(Y)$$

and the Cauchy (and dual Cauchy) identities

$$(2) \quad \sum_{\nu} s_{\nu/\lambda}(X) s_{\nu/\mu}(Y) = \prod_{i,j} \frac{1}{1-x_i y_j} \sum_{\kappa} s_{\lambda/\kappa}(Y) s_{\mu/\kappa}(X),$$

$$(3) \quad \sum_{\nu} s_{\nu/\lambda}(X) s_{\nu'/\mu'}(Y) = \prod_{i,j} (1+x_i y_j) \sum_{\kappa} s_{\lambda'/\kappa'}(Y) s_{\mu/\kappa}(X).$$

Schur functions can also be defined as generating series of semi-standard Young tableaux. I.e., for $\mu \subseteq \lambda$, $s_{\lambda/\mu}(x_1 \dots x_n) = \sum_T x_i^{\# \text{ of } i \text{ in } T}$ where the sum ranges over all semi-standard Young tableaux of shape λ/μ (fillings of the skew diagram λ/μ with numbers 1 through n such that the numbers weakly increase in rows left to right and strictly increase down columns). The fact the two definitions given so far are equivalent is the content of the Lindström-Gessel-Viennot lemma, see e.g. [16]. The Cauchy identities can be proven from this definition via the RSK correspondence [9].

In the infinite wedge formalism (we refer the reader to, e.g., [12, 13] for details), to each partition λ one associates a basis vector (respectively covector) $|\lambda\rangle$ (respectively $\langle\lambda|$) in the half-infinite wedge vector space denoted in the literature by $\bigwedge^{\infty} V$ (respectively in the dual vector space). We add the bilinear form defined by $\langle\lambda|\mu\rangle = \delta_{\lambda,\mu}$. It is useful to note that the identity operator acting on this space has the spectral decomposition $\sum_{\lambda} |\lambda\rangle\langle\lambda|$. Two important operators (each depending on a parameter) acting on these vectors are $\Gamma_-(x)$ and $\Gamma_+(y)$. We call them *half vertex operators*. One informally thinks of $\Gamma_-(x)$ as removing a horizontal strip from a partition (a strip weighted as $x^{\# \text{ boxes}}$) and $\Gamma_+(y)$ as adding a horizontal strip. More precisely:

$$\Gamma_+(z)|\lambda\rangle := \sum_{\mu \prec \lambda} z^{|\lambda|-|\mu|} |\mu\rangle,$$

$$\Gamma_-(z)|\lambda\rangle := \sum_{\mu \succ \lambda} z^{|\mu|-|\lambda|} |\mu\rangle.$$

We also define the following operators:

$$\tilde{\Gamma}_+(z) = \Gamma_+^{-1}(-z), \quad \tilde{\Gamma}_-(z) = \Gamma_-^{-1}(-z)$$

which we can think of as adding/removing vertical strips, i.e.:

$$\tilde{\Gamma}_+(z)|\lambda\rangle := \sum_{\mu \prec' \lambda} z^{|\lambda|-|\mu|} |\mu\rangle,$$

$$\tilde{\Gamma}_-(z)|\lambda\rangle := \sum_{\mu \succ' \lambda} z^{|\mu|-|\lambda|} |\mu\rangle.$$

These operators satisfy the following commutation relations (we remark the last three all follow from the first one):

$$(4) \quad \Gamma_+(y)\Gamma_-(x) = \frac{1}{1-xy} \Gamma_-(x)\Gamma_+(y),$$

$$(5) \quad \tilde{\Gamma}_+(y)\tilde{\Gamma}_-(x) = \frac{1}{1-xy} \tilde{\Gamma}_-(x)\tilde{\Gamma}_+(y),$$

$$(6) \quad \tilde{\Gamma}_+(y)\Gamma_-(x) = (1+xy)\Gamma_-(x)\tilde{\Gamma}_+(y),$$

$$(7) \quad \Gamma_+(y)\tilde{\Gamma}_-(x) = (1+xy)\tilde{\Gamma}_-(x)\Gamma_+(y).$$

Schur functions take the following form in terms of the Γ operators ($\mu \subseteq \lambda$):

$$s_{\lambda/\mu}(x_1 \dots x_n) = \langle\mu|\Gamma_+(x_1) \dots \Gamma_+(x_n)|\lambda\rangle$$

$$= \langle\lambda|\Gamma_-(x_1) \dots \Gamma_-(x_n)|\mu\rangle,$$

$$s_{\lambda'/\mu'}(x_1 \dots x_n) = \langle\mu|\tilde{\Gamma}_+(x_1) \dots \tilde{\Gamma}_+(x_n)|\lambda\rangle$$

$$= \langle\lambda|\tilde{\Gamma}_-(x_1) \dots \tilde{\Gamma}_-(x_n)|\mu\rangle.$$

In view of this, the commutation relations (4) and (6) are equivalent to the Cauchy (2) and dual Cauchy (3) identities (where now $X = (x), Y = (y)$ are one letter alphabets). The other two commutation relations are equivalent to the two mentioned by just conjugating all the partitions involved (e.g., (5) is really the 1-variable Cauchy identity (2) with all partitions primed). In Section 4 we will provide bijective proofs of (4) and (7) (that is, of the Cauchy and dual Cauchy identities).

3 Schur process and random tilings

The Schur process [12, 1] can be used to describe various statistical mechanical objects, including random plane and skew plane partitions [12, 13], the Aztec diamond [8], pyramid partitions [17], as well as the more general class of domino tilings called *steep tilings* [3]. In this note we will be concerned with exact sampling of all of the above from the q^{Volume} distribution, though we will keep generic parameters whenever issues of convergence can be ignored.

We start with a word $\gamma = \Gamma_1 \Gamma_2 \dots \Gamma_n$ where each $\Gamma \in \{\Gamma_-(y?), \Gamma_+(x?), \tilde{\Gamma}_-(y?), \tilde{\Gamma}_+(x?)\}$. We encode such a word in a (French style) Young diagram as depicted in Figure 3, and we call this the *encoded shape*. From a tiling perspective (see below) we are interested in two types of words. In one, there are no $\tilde{\Gamma}$'s, i.e., all $\Gamma_i \in \{\Gamma_-(y?), \Gamma_+(x?)\}$. This case leads to (skew) plane partitions. In another, the Γ_i 's alternate between $\Gamma_?$ and $\tilde{\Gamma}_?$, and this corresponds to domino tilings. Put a different way, the word $\Gamma_1 \Gamma_2 \dots \Gamma_n$ is really composed of two different words. One, $a_1 \dots a_n$, consists of booleans $a_i \in \{\text{add, remove}\}$ (this corresponds to whether Γ_i has a plus or minus sign as index and so to whether we are adding or removing a strip from a partition). The second, $b_1 \dots b_n$, consists of booleans $b_i \in \{\text{horizontal, vertical}\}$ (this corresponds to whether Γ_i is a Γ respectively $\tilde{\Gamma}$ and means we are adding/removing a horizontal respectively vertical strip). The a word gives the Young diagram that is the encoded shape: we move left to right, top to bottom; we start on the vertical axis and must end on the horizontal axis by drawing a horizontal line (left to right) for an add step and a vertical line (top to bottom) for a remove step. We call this encoded shape π and we represent it in French notation (see Figure 3).

For the word $\gamma = \Gamma_1 \Gamma_2 \dots \Gamma_n$, the Schur process of word γ with parameters $Z = (z_1, \dots, z_n)$ (z_i is either an x or a y parameter) is the measure on the set of sequences of partitions of the form $\Lambda = (0 = \lambda(0), \lambda(1), \dots, \lambda(n) = 0)$ given by

$$(8) \quad \text{Prob}(\Lambda) \propto \prod_{i=1}^n t(\Gamma_i, z_i, \lambda(i-1), \lambda(i))$$

where $t(\Gamma_i, z_i, \lambda(i-1), \lambda(i)) := \langle \lambda(i-1) | \Gamma_i(z_i) | \lambda(i) \rangle$ ($= z_i^{|\lambda(i)| - |\lambda(i-1)|}$ if $(\lambda(i), \lambda(i-1))$ appropriately interlace or dually interlace and $= 0$ otherwise) and z_i is one of the x parameters if $\Gamma_i \in \{\Gamma_+, \tilde{\Gamma}_+\}$ or one of the y parameters if $\Gamma_i \in \{\Gamma_-, \tilde{\Gamma}_-\}$. See Figure 3. That is, at each step in the partition process Λ , we either add or remove a vertical or horizontal strip, starting at 0 and ending at 0.

The most natural specialization of this measure (which also ensures convergence) is to first choose a parameter $0 < q < 1$ and then to choose the z parameters such that

$$\text{Prob}(\Lambda) \propto q^{\text{Volume}(\Lambda)} = q^{\sum_i |\lambda(i)|}.$$

This can be accomplished by the following:

$$z_i = \begin{cases} q^{-i} & \text{if } \Gamma_i \in \{\Gamma_+, \tilde{\Gamma}_+\}, \\ q^i & \text{if } \Gamma_i \in \{\Gamma_-, \tilde{\Gamma}_-\}. \end{cases}$$




Figure 3: The encoding of two Schur processes (read top-left to bottom-right), a skew plane partition (top) with $n = 6, \gamma = \Gamma_+ \Gamma_+ \Gamma_- \Gamma_- \Gamma_+ \Gamma_-$, $Z = (x_1, x_2, y_3, x_3, y_2, y_1)$ and a steep domino tiling (bottom) with $n = 8, \gamma = \Gamma_+ \tilde{\Gamma}_+ \Gamma_- \tilde{\Gamma}_- \Gamma_+ \tilde{\Gamma}_+ \Gamma_- \tilde{\Gamma}_-$, $Z = (x_1, x_2, y_4, y_3, x_3, y_2, x_4, y_1)$. The encoding Young diagrams are $(3, 3, 2)$ (top) and $(4, 3, 2, 2)$ (bottom). Each bullet \bullet represents a partition (and the process/sequence of partitions Λ is what one obtains reading the bullets clockwise from the vertical axis to the horizontal axis). Asterisks $*$ also represent partitions, important in Section 4.

There are several special cases of interest, listed below. They all correspond to various types of tilings. We also list the q^{Volume} specialization in the x, y parameter setting (where every add step gets an x parameter and every remove step gets a y parameter).

1) Plane partitions with base in an $m \times n$ rectangle and distributed as q^{Volume} . $\gamma = \Gamma_+ \dots \Gamma_+ \Gamma_- \dots \Gamma_-$ where there are m Γ_+ 's and n Γ_- 's, $Z = (x_1, \dots, x_m, y_n, \dots, y_1)$ and $x_i = q^{m-i}, y_i = q^{n-i}$. The encoded shape is an $m \times n$ rectangle.

2) Skew plane partitions. We concatenate together several words γ and several parameter sets Z of the form given in the above example. Due to cumbersome notation, we only give a simple (but yet nontrivial) example. Let m_1, m_2, n_1, n_2 be strictly

positive integers with $m_1 + m_2 = m, n_1 + n_2 = n$. Let $\gamma = \Gamma_+ \dots \Gamma_+ \Gamma_- \dots \Gamma_- \Gamma_+ \dots \Gamma_+ \Gamma_- \dots \Gamma_-$ (there are m_1 Γ_+ 's in the first/leftmost cluster and m_2 in the second; there are n_2 Γ_- 's in the first cluster and n_1 in the second), $Z = (x_1, \dots, x_{m_1}, y_n, \dots, y_{n_1+1}, x_{m_1+1}, \dots, x_m, y_{n_1}, \dots, y_1)$,
 $x_i = q^{m_1-i} \delta_{1 \leq i \leq m_1} + q^{m_1-n_2-i} \delta_{m_1+1 \leq i \leq m}$,
 $y_i = q^{m_2+n-i} \delta_{1 \leq i \leq n_1} + q^{n-i} \delta_{n_1+1 \leq i \leq n}$. An example of an encoding is given in Figure 3 ($m_1 = n_1 = 2, m_2 = n_2 = 1$). The encoded shape is (generically) a Young diagram contained in the $m \times n$ rectangle where m is the number of x 's and n is the number of y 's.

3) $n \times n$ Aztec diamonds distributed as q^{Volume} . $\gamma = \tilde{\Gamma}_+ \Gamma_- \tilde{\Gamma}_+ \Gamma_- \dots \tilde{\Gamma}_+ \Gamma_-$ (there are a total of $2n$ Γ 's), $Z = (x_1, y_n, x_2, y_{n-1}, \dots, x_n, y_1)$ and $x_i = q^{-2i}, y_i = q^{2n-2i}$. The encoding shape of an Aztec diamond is a staircase partition. Also notice in this case the geometry forces all partitions λ belonging to the process to be bounded inside the n^n square, so there are no convergence issues and one can get any general Schur distribution with generic x and y parameters.

4) Pyramid partitions contained in an $m \times n$ rectangle. Assume (for simplicity) m, n are even. $\gamma = \tilde{\Gamma}_+ \Gamma_+ \dots \tilde{\Gamma}_+ \Gamma_+ \tilde{\Gamma}_- \Gamma_- \dots \tilde{\Gamma}_- \Gamma_-$ (there are m Γ 's indexed by a plus sign and n by a minus sign), $Z = (x_1, \dots, x_m, y_n, \dots, y_1)$ and $x_i = q^{m-i}, y_i = q^{n-i}$. The encoded shape is an $m \times n$ rectangle.

5) Steep tilings. We concatenate together several words γ and several parameter sets Z of the form given in the above example. Again due to cumbersome notation, we only provide a simple yet non-trivial example. Let m_1, m_2, n_1, n_2 be strictly positive (and for convenience) even integers with $m_1 + m_2 = m, n_1 + n_2 = n$. We have $\gamma = \Gamma_+ \tilde{\Gamma}_+ \dots \Gamma_+ \tilde{\Gamma}_+ \Gamma_- \tilde{\Gamma}_- \dots \tilde{\Gamma}_- \Gamma_+ \dots \tilde{\Gamma}_+ \Gamma_- \dots \tilde{\Gamma}_-$ (there are m_1 Γ 's indexed by a plus sign in the first/leftmost cluster and a total of m ; there are n_2 Γ 's indexed by a minus sign in the leftmost cluster and a total of n). Z can be taken the same as in the case of skew plane partitions, and the rest of the details of that example also apply verbatim. An example of an encoding (this time more general, with $m_1 = n_3 = 2, m_2 = m_3 = n_1 = n_2 = 1$) is given in Figure 3. The encoded shape is (generically) a Young diagram contained in a $m \times n$ rectangle where m is the number of x 's and n is the number of y 's.

Plane partitions (and more generally skew plane partitions) correspond to lozenge tilings - see Figure 4; we read the partition process by reading the

heights of the horizontal lozenges on each vertical slice bisecting such lozenges, left to right. Aztec diamonds, pyramid partitions and more generally steep tilings correspond to domino tilings - see Figures 4 and 5 and we use Maya diagrams to encode the par-




Figure 4: A skew plane partition corresponding to interlacing sequence $0 \prec (1) \prec (3, 1) \prec (4, 2) \succ (2, 2) \succ (2) \prec (3, 2) \prec (4, 2) \succ (2) \succ 0$ (partitions are given by heights of horizontal lozenges on vertical slices, left to right) and a (full) 2×2 Aztec diamond corresponding to $0 \prec' (1, 1) \succ (1) \prec' (2) \succ 0$ (Maya diagrams for partitions given on blue diagonals as indicated).




Figure 5: A pyramid partition (left) corresponding to the sequence $0 \prec' (1) \prec (1, 1) \prec' (2, 2) \prec (2, 2, 2) \prec' (3, 3, 2) \succ (3, 2) \succ' (2, 1) \succ (2) \succ' (1) \succ 0$ and a more general steep domino tiling (right). The partitions are represented diagonally (the blue lines and lines parallel to them) as Maya diagrams, read in a SE-NW direction.

4 Bijective sampling of Schur processes

Perfect sampling algorithms based on partition processes for classes of tilings including plane and skew plane partitions [1] and Aztec diamonds (see [5] and remark below) have been proposed before. There is also the coupling from the past approach of Propp and Wilson [15]. In this section we give an overall encompassing algorithm for the Schur process (different from, but similar to, that of [1]) which under appropriate specializations gives perfect random sampling of plane (and skew plane) partitions, Aztec diamonds, pyramid partitions and more generally, steep tilings. The algorithm is based on bijective proofs of the Cauchy identities for Schur functions.

We now give bijective proofs of the Cauchy identity (2) and the dual Cauchy identity (3) when $X = (x)$ and $Y = (y)$ are single letter alphabets (the general case is then proved by induction using the branching rule (1)). We refer the reader to the works of Gessel [6], Pak–Postnikov [14] and Krattenthaler [10] for related work on the subject. Alternatively, we give bijective proofs of the commutation relation between $\Gamma_+(x)$ and $\Gamma_-(y)$, respectively $\tilde{\Gamma}_+(x)$ and $\tilde{\Gamma}_-(y)$. We represent this schematically in Figure 6.

Cauchy case. Let λ, κ, μ be three partitions such that $\lambda \succ \kappa, \mu \succ \kappa$ and let $G \in \mathbb{N}$. We describe a procedure (which appeared first in Gessel [6]) for building a fourth partition ν with the properties that $\lambda \prec \nu, \mu \prec \nu$ and $|\lambda| + |\mu| + G = |\kappa| + |\nu|$, in such a way that the mapping $(\kappa, G) \mapsto \nu$ is bijective (i.e. every possible ν is obtained once and exactly once by the procedure). See Figure 6 (top) for a schematic representation. We construct ν (which has at most $\max(\ell(\lambda), \ell(\mu)) + 1$ parts) by setting

$$\nu_i = \begin{cases} \max(\lambda_1, \mu_1) + G & \text{if } i = 1 \\ \max(\lambda_i, \mu_i) + \min(\lambda_{i-1}, \mu_{i-1}) - \kappa_{i-1} & \text{if } i > 1 \end{cases}$$

and it is readily checked that the wanted properties hold. We deduce the identity

$$\sum_{\substack{\nu \\ \nu \succ \lambda \\ \nu \succ \mu}} x^{|\nu| - |\lambda|} y^{|\nu| - |\mu|} = \frac{1}{1 - xy} \sum_{\substack{\kappa \\ \kappa \prec \lambda \\ \kappa \prec \mu}} x^{|\mu| - |\kappa|} y^{|\lambda| - |\kappa|}$$

which amounts to the commutation relation (4), and to the Cauchy identity (2) in the case where X and Y are reduced to a single variable.

To obtain a bijective proof of the relation (5), one applies the same procedure after apriori conjugating all the partitions (and at the end, conjugating the resulting ν again).

The first bijective procedure (proving (4)) has time complexity $O(\ell(\lambda))$ while the second (proving (5)) is $O(\max(\ell(\lambda), \lambda_1))$ so both are commonly $O(\max(\ell(\lambda), \lambda_1))$ (and we will need this weaker bound in order to give uniform estimates below).

The above bijection can be turned into a random sampling procedure easily. With the notation set up above, define a method `sampleHH` (HH stands for horizontal-horizontal and the fact that we are commuting Γ_+ and Γ_-) that does the following:

```
def sampleHH( $\lambda, \mu, \kappa, x, y$ )
  sample  $G \sim \text{Geom}(xy)$ 
  construct  $\nu$  based on the bijective procedure described above
```

```
return  $\nu$ 
```

where $\text{Geom}(\xi)$ samples a single geometric random variable from the distribution $\text{Prob}(k) \propto \xi^k$ (we assume $0 \leq \xi < 1$). Then applying this procedure will produce ν distributed as $\text{Prob}(\nu) \propto s_{\nu/\lambda}(x)s_{\nu/\mu}(y)$ using minimal entropy (the sampling of a single uniform geometric variable) and assuming λ, ν, κ are coming from a Schur process, in the end so will ν (from a different Schur process, of course).

A method `sampleVV` can be defined in the same way and corresponds to commuting $\tilde{\Gamma}_+$ and $\tilde{\Gamma}_-$.

Dual Cauchy case. Let λ, κ, μ be three partitions such that $\kappa \prec' \lambda, \mu \succ \kappa$ and let $B \in \{0, 1\}$. We describe a procedure similar to the above one for building a fourth partition ν with the properties that $\lambda \prec \nu, \mu \prec' \nu$ and $|\lambda| + |\mu| + B = |\kappa| + |\nu|$, in such a way that the mapping $(\kappa, B) \mapsto \nu$ is bijective. See Figure 6 (bottom) for a schematic representation. Here we directly define the method `sampleHV` (HV stands for horizontal-vertical and the fact that we are commuting Γ_+ and $\tilde{\Gamma}_-$):

```
def sampleHV( $\lambda, \mu, \kappa, x, y$ )
  sample  $B \sim \text{Bernoulli}(\frac{xy}{1+xy})$ 
  for  $i = 1 \dots \max(\ell(\lambda), \ell(\mu)) + 1$ 
    if  $\lambda_i \leq \mu_i < \lambda_{i-1}$  then  $\nu_i = \max(\lambda_i, \mu_i) + B$ 
    else  $\nu_i = \max(\lambda_i, \mu_i)$ 
    if  $\mu_{i+1} < \lambda_i \leq \mu_i$  then  $B = \min(\lambda_i, \mu_i) - \kappa_i$ 
  return  $\nu$ 
```

where $\text{Bernoulli}(\xi)$ is a Bernoulli random variable that returns 0 with probability $1 - \xi$ and 1 with probability ξ . To check the validity of our method, note first that the interlacing conditions for κ and ν amount to

$$\begin{aligned} \max(\lambda_i - 1, \mu_{i+1}) &\leq \kappa_i \leq \min(\lambda_i, \mu_i), \\ \max(\lambda_i, \mu_i) &\leq \nu_i \leq \min(\lambda_{i-1}, \mu_i + 1) \end{aligned}$$

where by convention $\lambda_0 = \infty$. In particular, the quantity $\min(\lambda_i, \mu_i) - \kappa_i$ vanishes unless

$$(9) \quad \mu_{i+1} < \lambda_i \leq \mu_i$$

in which case it may also take the value 1. Let $i_1 < i_2 < \dots < i_r$ be the i 's such that (9) holds. Similarly, the quantity $\nu_i - \max(\lambda_i, \mu_i)$ vanishes unless

$$(10) \quad \lambda_i \leq \mu_i < \lambda_{i-1}$$

in which case it may also take the value 1. Let $i'_1 < i'_2 < \dots < i'_s$ be the i 's such that (10) holds. Our method works provided that $s = r + 1$ and

$$i'_1 \leq i_1 < i'_2 \leq i_2 < \dots < i'_r \leq i_r < i'_{r+1},$$

which follows from the fact that the mapping

$$i \mapsto \min\{j > i, \lambda_j \leq \mu_i\}$$

defines a bijection between $\{0, i_1, i_2, \dots, i_r\}$ and $\{i'_1, i'_2, \dots, i'_s\}$.

The procedure bijectively proves that

$$\sum_{\substack{\nu \succ \lambda \\ \mu \prec' \nu}} x^{|\nu| - |\lambda|} y^{|\nu| - |\mu|} = (1 + xy) \sum_{\substack{\kappa \prec' \lambda \\ \kappa \prec \mu}} x^{|\mu| - |\kappa|} y^{|\lambda| - |\kappa|}$$

which amounts to the commutation relation (7), and to the dual Cauchy identity (3) in the case where X and Y are reduced to a single variable. As in the sampleHH case, the complexity is $O(\ell(\lambda))$ which as before is bounded above by $O(\max(\ell(\lambda), \lambda_1))$ (recall we will need this weaker bound to produce uniform estimates below). Applying the procedure will produce ν distributed as $Prob(\nu) \propto s_{\nu/\lambda}(x) s_{\nu'/\mu'}(y)$ using minimal entropy (the sampling of a single Bernoulli random variable) and assuming λ, ν, κ are coming from a Schur process, in the end so will ν (from a different Schur process, of course).

A method sampleVH, of the same complexity, is defined by exchanging the roles of λ and μ and corresponds to commuting $\tilde{\Gamma}_+$ and Γ_- according to (6).




Figure 6: A diagrammatic representation of the two bijections and of the atomic sampling steps corresponding to commuting $\Gamma_+(x)$ and $\Gamma_-(y)$ (top), respectively $\Gamma_+(x)$ and $\tilde{\Gamma}_-(y)$ (bottom). $G \in \mathbb{N}, B \in \{0, 1\}$.

We now give the exact sampling algorithm for a Schur process with word $\gamma = \Gamma_1 \dots$ and parameter list $Z = (x_1, \dots, y_1)$. Suppose we have m x parameters and n y parameters. First, (as a pre-computation step) from γ we build the encoding shape (like in Figure 3) which we call π . We also build a function that from γ returns for each $(i, j) \in \pi$ one of HH, VV, HV or VH based on whether the two Γ operators (with parameters

x_i and y_j) are $(\Gamma_+, \Gamma_-), (\tilde{\Gamma}_+, \tilde{\Gamma}_-), (\Gamma_+, \tilde{\Gamma}_-)$ or $(\tilde{\Gamma}_+, \Gamma_-)$ respectively. Let us denote this method by $getType((i, j) \in \pi)$. We can wrap the four atomic sample steps described above in a single method which we call sample:

```
def sample( $\lambda, \mu, \kappa, x, y, type$ )
  case type:
    HH: return sampleHH( $\lambda, \mu, \kappa, x, y$ )
    HV: return sampleHV( $\lambda, \mu, \kappa, x, y$ )
    VH: return sampleVH( $\lambda, \mu, \kappa, x, y$ )
    VV: return sampleVV( $\lambda, \mu, \kappa, x, y$ )
```

The idea of the algorithm is to build the shape π (see Figure 3) one square at a time, starting from the empty partitions on the coordinate axes and at the end reading the partitions corresponding to the dots on the boundary of π that lie in the first quadrant (the dots in black in Figure 3 and the two in red on the two axes). Note below $\ell(\pi) = \pi'_1 = n, \pi_1 = m$.

Algorithm SchurSample

Input: π , partitions $\tau(0, i) = 0 = \tau(j, 0), 0 \leq i, j \leq \pi_1, \ell(\pi)$
for $j = 1 \dots \ell(\pi)$

for $i = 1 \dots \pi_j$
 $type = getType(i, j)$
 $\tau(i, j) = sample(\tau(i-1, j), \tau(i, j-1), \tau(i-1, j-1), x_i, y_j, type)$

Output: The sequence of partitions Λ defined by

$$\Lambda := (0 = \tau(l_0), \tau(l_1), \dots, \tau(l_{m+n-1}), \tau(l_{m+n}) = 0),$$

where

$$L = (l_0 = (0, \pi'_1), l_1 = (1, \pi'_1), \dots, (\pi_1, 1), l_{m+n} = (\pi_1, 0))$$

is the ordered sequence of lattice points on the boundary of π clockwise from the vertical axis to the horizontal axis.

Remark. The final result of the sampling does not depend on the ordering in which we sample the inner corner partitions $\tau(i, j)$ as long as we respect the causality relations: $\tau(i_1, j_1)$ needs to be sampled before $\tau(i_2, j_2)$ if $i_2 \geq i_1$ and $j_2 \geq j_1$.

Theorem 1. *The algorithm SchurSample produces an exact random sample from the Schur process corresponding to word γ and parameter list Z .*

Proof. The proof is by induction on $|\pi|$. If $|\pi| = 1$ ($\pi = (1)$) this is immediate. For example, if $\gamma = \Gamma_+(x)\Gamma_-(y)$ then we are sampling ν based on $\lambda = \mu = \kappa = 0$ using the sampleHH procedure, which produces the right result (assuming one can exactly sample a geometric random variable). That is, it produces ν distributed as $Prob(\nu) \propto s_\nu(x)s_\nu(y)$. The reasoning is similar for the other three types of possible words.

Suppose we can perfectly sample any Schur pro-

cesses encoded by π for any π of weight less than or equal to k . Take a Schur process Λ with $|\pi| = k + 1$ and word $\gamma = \Gamma_1 \dots \Gamma_m$. Remove a (any) outer corner from π and call the result π^0 . Note $|\pi^0| = |\pi| - 1 = k$. The result encodes a Schur process Λ^0 with word γ^0 which is obtained from γ by interchanging (in γ) some $\Gamma_i(x)$ ($\Gamma_i \in \{\Gamma_+, \tilde{\Gamma}_+\}$) with $\Gamma_{i+1}(y)$ ($\Gamma_{i+1} \in \{\Gamma_-, \tilde{\Gamma}_-\}$). By the inductive hypothesis we can sample this “smaller” process exactly and correctly. We then perform an extra atomic step to add the missing corner to π^0 to obtain π . That is, if λ, κ, μ are the three partitions that sit at the inner corner of π^0 (like in Figure 6), we sample the outer corner ν based on the three using one of the four sample{HH,HV,VH,VV} procedures (chosen based on Γ_i and Γ_{i+1} which need to be interchanged). The atomic step exactly samples this corner so that it fits into the new process correctly (everything depends only on the 3 partitions involved and the type of atomic step). Thus we obtain Λ sampled exactly and correctly from Λ^0 by replacing (in Λ^0) the partition κ with the partition ν . \square

Remark. For the Aztec diamond, the algorithm `SchurSample` is related to the domino shuffling algorithm of [5]. More precisely, we leave as an exercise for the reader to check that, applying the growth procedure to the diagram π along anti-diagonals in the case of the Aztec diamond exactly gives the domino shuffling algorithm (hint: this is a simple check, once one expands the definitions of all the combinatorial constructions present in this statement).

Remark. As is clear from the construction, any anti-diagonal path (going right and down) from the vertical axis to the horizontal axis that appears in the construction of the Schur process encoded by π (i.e., a subpartition of π in Figure 3) will itself encode a Schur process. However, it will seldom be related in any meaningful way with the process encoded by π (for example, while π may correspond to a domino tiling, a subpartition of it may not).

Remark. One can modify the algorithm slightly for better space complexity (but the same time complexity). With $m = \pi_1, n = \ell(\pi)$, one can start with the $m + n + 1$ red dots depicting zero partitions in Figure 3 and sample using the local rules, but update partitions in place: that is, for example $\tau(1,1)$ would be sampled and then overwritten in the place of $\tau(0,0)$ since the latter is not needed

anymore. Similarly $\tau(1,2)$ would be sampled and then overwritten in the place of $\tau(0,1)$ using the same logic, and so on.

The explicit description of the algorithm immediately allows us to estimate its time complexity, which is random and depends on the output. We assume that we can sample Bernoulli and geometric random variables in time $O(1)$. We have:

Proposition 1. *The time complexity of `SchurSample` is $O(|\pi|L)$, where we recall that π is the encoding shape of the Schur process, and where $L := \max\{\tau(l_i)_1, \ell(\tau(l_i)), i \in [0 \dots m + n]\}$ is the maximum of the maximum part and the maximum number of parts of the partitions in the output sequence Λ .*

Proof. This is clear by induction on $|\pi|$. Indeed, we have seen that the complexity of adding a box to π (Cauchy case and dual Cauchy case) is $O(\max(\ell(\lambda), \lambda_1))$ where λ is the top-right partition in the added box. The result thus follows by observing that the quantity L cannot decrease when we add a box to π , which is clear by examining the rules. \square

Remark. In some cases we can deterministically bound the quantity L above, thus giving a better (deterministic) bound on the complexity. Suppose we sampled a Schur process with word γ and shape π , and suppose further that in doing so we never had to use both the `sampleHH` and the `sampleVV` procedure (this fact of course only depends on the word γ). For simplicity, suppose we never had to use the `sampleVV` procedure (we may have used the `sampleHH` procedure multiple times though). Then under these assumptions, the time complexity of `SchurSample` is $O(|\pi|A)$ where A is the total number of add steps (of Γ_+ and $\tilde{\Gamma}_+$ in γ).

We now address entropy minimality/optimalty. Consider a Schur process with word γ and encoded shape π . The algorithm above samples from it exactly and bijectively. In particular, it constructs an array (partial matrix of shape π) $(S_{ij})_{(i,j) \in \pi}$ of $|\pi|$ positive integers (some of which are in $\{0, 1\}$, while others are unrestricted). This array is in bijection with the Schur process via the local bijections obtained to sample said process. Thus because of the bijection, we need exactly $|\pi|$ positive integers (with appropriate restrictions) to sample such a process, and our algorithm samples using this exact amount of randomness (we note that this translates into sampling $|\pi|$ pseudo-random uniform $(0, 1)$ variables in the implementation we provided). We have thus proven:

Proposition 2. *Assuming that we can sample sequences of independent geometric random variables and random bits optimally, the algorithm SchurSample is entropy optimal.*

Remark. In the case of plane partitions with base inside an $m \times n$ rectangle (encoded shape $\pi = m^n$), the bijection is given by the RSK correspondence (see [6]). It takes an $m \times n$ matrix with positive integers (the one obtained from running the algorithm) and maps it to a pair of semi-standard Young tableaux of common shape λ . One then takes these 2 tableaux and conjoins them (view each as a sequence of interlacing partitions starting from 0 and ending in λ) to obtain a plane partition (with central slice given by λ) which is the sampled plane partition.

Remark. The algorithm above can be thus viewed as a generalized RSK correspondence. It takes a shape π and a partial matrix $(S_{ij})_{(i,j) \in \pi}$ of (geometric/Bernoulli) integers and a word γ (from which π was built) and produces a sequence of partitions which interlace or dually interlace at every step. It reduces to RSK as remarked above (but also to dual RSK). See also [6, 10, 14].

Below we present some samples from the algorithm. In Figure 7 we present a large plane partition (top) and Aztec diamond (bottom). The deterministic asymptotic shape is a law of large numbers (see [12] for the case of the plane partition and [7] for the Aztec diamond). In Figure 8, we show a large pyramid partition. Its width and the parameter q are tuned so that its apparent limit shape exhibits two interesting cusp points. This will be studied theoretically in future work.

5 Conclusion

In this article we have presented an efficient polynomial time algorithm for sampling from any Schur process. It makes minimal use of randomness and leads to efficient sampling algorithms for a variety of tilings including plane partitions, the Aztec diamond and pyramid partitions. One can speculate whether the Schur functions can be replaced by more general symmetric functions. The answer is yes, but in order to preserve polynomiality, the branching (Pieri) coefficients for such functions should not depend on the skew diagram. For example, while one can extend the formalism described in Section 3 to Macdonald processes [2], it is unlikely an efficient sampling algorithm will exist. However, Macdonald measures (or more degenerate versions that still generalize Schur)




Figure 7: A random large plane partition with base contained in an 100×100 box and $q = 0.93$ (top) and a random 100×100 Aztec diamond (bottom) with $q = 0.99$. Both exhibit deterministic limit shapes. The limit shape is known in the second case as the “arctic circle” (the terminology comes from the uniform case $q = 1$, in which case the arctic circle... is indeed a circle).

on steep tilings are of interest independent of sampling, and we hope to address this in future publications. Also of interest is the case where one replaces Schur by Schur’s P -polynomials (a degenerate case of the Macdonald hierarchy). Here, because of the “free-fermionic” nature of the problem, there may be hope for a fast perfect sampling algorithm.

In a different direction, random sampling can be used to conjecture and prove various laws of large numbers for Schur processes. Some such laws have been proven, but more await discovery, especially




Figure 8: A random pyramid partition of width 100 with parameter $q = .99$ (only the particles of the corresponding Maya diagrams are displayed, not the dominos). The apparent limit shape (in green) seems to exhibit cusp points (in blue), an interesting phenomenon that will be the subject of future work.

regarding steep tilings, as we illustrated on Figure 8. This will be addressed in future work.

Acknowledgments

The authors acknowledge financial support from the Agence Nationale de la Recherche via the grants ANR-08-JCJC-0011 “IComb” (SC), ANR-10-BLAN-0123 “MAC2” (DB, CB) and ANR-12-JS02-0001 “Cartaplus” (JB, GC), and from the Ville de Paris via Projet Émergences “Combinatoire à Paris” (JB, GC, SC).

References

- [1] A. BORODIN, *Schur dynamics of the Schur processes*, Adv. Math., 228 (2011), pp. 2268–2291.
- [2] A. BORODIN AND I. CORWIN, *Macdonald processes*, Probab. Theory Related Fields, 158 (2014), pp. 225–400.
- [3] J. BOUTTIER, G. CHAPUY, AND S. CORTEEL, *From Aztec diamonds to pyramids: steep tilings*, arXiv:1407.0665, (2014).
- [4] S. DESREUX AND E. RÉMILA, *An optimal algorithm to generate tilings*, J. Discrete Algorithms, 4 (2006), pp. 168–180.
- [5] N. ELKIES, G. KUPERBERG, M. LARSEN, AND J. PROPP, *Alternating-sign matrices and domino tilings. I*, J. Algebraic Combin., 1 (1992), pp. 111–132.
- [6] I. M. GESSEL, *Counting paths in Young’s lattice*, J. Statist. Plann. Inference, 34 (1993), pp. 125–134.
- [7] W. JOCKUSCH, J. PROPP, AND P. SHOR, *Random domino tilings and the arctic circle theorem*, arXiv:math/9801068, (1998).
- [8] K. JOHANSSON, *The arctic circle boundary and the Airy process*, Ann. Probab., 33 (2005), pp. 1–30.
- [9] D. E. KNUTH, *Permutations, matrices, and generalized Young tableaux*, Pacific J. Math., 34 (1970), pp. 709–727.
- [10] C. KRATTENTHALER, *Growth diagrams, and increasing and decreasing chains in fillings of Ferrers shapes*, Adv. in Appl. Math., 37 (2006), pp. 404–431.
- [11] I. G. MACDONALD, *Symmetric functions and Hall polynomials*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, second ed., 1995. With contributions by A. Zelevinsky, Oxford Science Publications.
- [12] A. OKOUNKOV AND N. RESHETIKHIN, *Correlation function of Schur process with application to local geometry of a random 3-dimensional Young diagram*, J. Amer. Math. Soc., 16 (2003), pp. 581–603 (electronic).
- [13] ———, *Random skew plane partitions and the Pearcey process*, Comm. Math. Phys., 269 (2007), pp. 571–609.
- [14] I. PAK AND A. POSTNIKOV, *Oscillating tableaux, $(S_p \times S_p)$ -modules, and Robinson-Schensted-Knuth correspondence*, in Proceedings of FPSAC ’96, Minneapolis, MN, 1996.
- [15] J. PROPP AND D. WILSON, *Coupling from the past: a user’s guide*, in Microsurveys in discrete probability (Princeton, NJ, 1997), vol. 41 of DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Amer. Math. Soc., Providence, RI, 1998, pp. 181–192.
- [16] R. P. STANLEY, *Enumerative combinatorics. Vol. 2*, vol. 62 of Cambridge Studies in Advanced Mathematics, Cambridge University

Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

- [17] B. YOUNG, *Generating functions for colored 3D Young diagrams and the Donaldson-Thomas invariants of orbifolds*, *Duke Math. J.*, 152 (2010), pp. 115–153. With an appendix by Jim Bryan.