



HAL
open science

Exact Random Generation of Symmetric and Quasi-symmetric Alternating-sign Matrices

Philippe Duchon, Florent Le Gac

► **To cite this version:**

Philippe Duchon, Florent Le Gac. Exact Random Generation of Symmetric and Quasi-symmetric Alternating-sign Matrices. GASCom, Jun 2008, Bibbiena, Italy. pp.131-141. hal-00412149

HAL Id: hal-00412149

<https://hal.science/hal-00412149>

Submitted on 1 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exact Random Generation of Symmetric and Quasi-symmetric Alternating-sign Matrices

Philippe Duchon and Florent Le Gac
LaBRI, ENSEIRB - Université Bordeaux 1

January 29, 2008

Abstract

We show how to adapt the Monotone Coupling from the Past exact sampling algorithm to sample from some symmetric subsets of finite distributive lattices. The method is applied to generate uniform random elements of all symmetry classes of alternating-sign matrices.

1 Introduction

Coupling from the Past (CFTP) is an algorithm due to Propp and Wilson [6] to sample from the exact stationary distribution (rather than from an approximation, like the classical MCMC method) of an ergodic Markov chain. In its arguably most practical version of MONOTONE-CFTP, the Markov chain's state space is a partially ordered set with unique minimum and maximum elements, and the chain is run through a *monotone update mechanism*: at each time step, an appropriately-distributed *increasing* function mapping the state space to itself is randomly chosen, and arbitrarily many copies of the Markov chain can be coupled using the same function by defining the new state of each of them to be the image of their previous state. The key element in the CFTP algorithm lies in being able to (backward) compose a number of such random update functions until the resulting function is found to be a constant. Using only *increasing* functions makes it possible to detect this coalescence by computing the successive images of only the minimum and maximum elements of the state space.

Thus, CFTP is particularly well suited to sampling from *distributive lattices*, either uniformly or according to some well-behaved distribution. In favorable cases, this makes it possible to completely eliminate initiation bias from the corresponding MCMC simulations, in time comparable to what a well-tuned (for small enough bias) MCMC algorithm would require. It is also worth mentioning that tight bounds on the mixing time of Markov chains are often tricky to obtain, so that MCMC simulations typically either have to be run for an empirically determined time, or for a sometimes widely overestimated guaranteed mixing time.

Many combinatorially interesting families of objects fall in the “monotone CFTP on a distributive lattice” framework, but in some cases this fails because of some imposed symmetry condition. As an example, consider *alternating-sign matrices* (ASM for short; see Section 3 for definitions). These are in bijection with another set of matrices which form a distributive lattice, and CFTP is an efficient way to generate random ASMs.

The group of symmetries of the unit square acts naturally on alternating-sign matrices, and for each of its subgroups, one can consider the class of matrices which are invariant under its action. Enumerative formulae for many symmetry classes were conjectured by Robbins [9], and proved by Kuperberg and others [4, 7, 8].

Random generation of non-symmetric ASMs can be considered a textbook exercise in MONOTONE-CFTP. For some symmetry classes, however, the defining symmetries correspond to a *decreasing* automorphism on the lattice of ASMs, and as a result the set of symmetric ASMs no longer has a natural ordering.

In this paper, we show how to adapt the monotone-CFTP algorithm in a systematic way to this kind of situation. The new variant, which we call SYMMETRIC-CFTP, uses the fact that the set we are trying to sample from can be seen as the set of symmetric ideals of some partially ordered set.

The paper is organized as follows. In Section 2, we review the usual MONOTONE-CFTP and describe the general idea of SYMMETRIC-CFTP. In Section 3, we apply it to random sampling from arbitrary symmetry classes of alternating-sign matrices, and give some experimental results.

2 Monotone and symmetric coupling from the past

Let (P, \leq) denote a finite partially ordered set (poset). Recall that a subset $E \subset P$ is an *ideal* of P if it is closed under taking smaller elements ($x \leq y$ and $y \in E$ imply $x \in E$), and that y is said to *cover* x if $x < y$ but no z exists satisfying $x < z < y$. The set $J(P)$ of lower ideals of P forms a finite distributive lattice, and it is a standard result (see, *e.g.*, [11]) that any finite distributive lattice is isomorphic to the set of lower ideals of some finite poset. We consider the problem of sampling uniformly at random from $J(P)$, or from a specific subset.

2.1 Monotone CFTP

The standard application of MONOTONE-CFTP in this context is as follows. For any $x \in P$, let T_x^+ and T_x^- denote the applications $J(P) \rightarrow J(P)$ defined by

$$T_x^+(E) = \begin{cases} E \cup \{x\} & \text{if } E \cup \{x\} \in J(P) \\ E & \text{otherwise} \end{cases}, T_x^-(E) = \begin{cases} E - \{x\} & \text{if } E - \{x\} \in J(P) \\ E & \text{otherwise} \end{cases}$$

Thus, T_x^+ (resp. T_x^-) acts on any ideal by trying to add (resp. remove) x , leaving it unchanged if this violates the ideal condition. Note that each such application is *monotone increasing* for the order inclusion, *i.e.* $E \subseteq F \Rightarrow T(E) \subseteq T(F)$.

Now consider any probability distribution π on P (subject to the condition that $\pi(x) > 0$ for all $x \in P$). We can define a Markov chain with state space $J(P)$ where one step of the chain consists in choosing a random $x \in P$ according to π and an independent sign ϵ uniformly in $\{+, -\}$, and replacing the current state E by $E' = T_x^\epsilon(E)$. Such a Markov chain is clearly irreducible (one can move from any state E to the empty ideal \emptyset by applying the transforms T_x^- , $x \in E$, in decreasing order, and back by applying the transforms T_x^+ , $x \in E$, in increasing order) and aperiodic ($T_x^-(E) = E$ whenever x is not a maximal element of E and $T_x^+(E) = E$ whenever x is not a minimal element of $P - E$, which implies that the transition matrix has nonzero diagonal elements), with a symmetric transition matrix; thus, its unique stationary distribution is the uniform distribution on the whole lattice $J(P)$. In fact, the above description actually defines a *grand coupling* of an arbitrary number of copies of the Markov chains: one can start one copy of the chain from each state and run them in parallel using the same update functions T_x^ϵ ; if at any time two copies are in the same state, they will be in the same state forever. With probability 1, all copies will ultimately be in the same state; this is equivalent to the fact that since one can obtain a constant function by composing the update functions T_x^ϵ (for instance, for any linear extension $x_1 < x_2 < \dots < x_n$ of the order \leq on P , the compose

$$T_{x_1}^- \circ T_{x_2}^- \circ \dots \circ T_{x_n}^-$$

is the constant function which maps each ideal to the empty ideal), composing random functions will ultimately yield a constant function.

The idea of CFTP is to compose such update functions in *reverse order* until the compose is constant, *i.e.* generate a sequence of independent pairs (x_i, ϵ_i) until $T_{1,n} = T_{x_1}^{\epsilon_1} \circ \dots \circ T_{x_k}^{\epsilon_k}$ is constant, and then output the single value of this function, which is uniformly distributed on $J(P)$. The *monotonicity* of the update functions is what makes the method practically useful: since each $T_{1,n}$ is an increasing mapping of $J(P)$ to itself, one can decide whether it is constant by computing the images of the minimum and maximum ideals \emptyset and P ; $T_{1,n}$ is constant if and only if $T_{1,n}(\emptyset) = T_{1,n}(P)$.

2.2 Automorphisms of finite posets

Now consider the group of *increasing or decreasing* automorphisms of P . *Increasing* automorphisms are the “true” automorphisms of the poset, that is, bijections σ of P to itself such that $\sigma(x) \leq \sigma(y) \iff x \leq y$; *decreasing* automorphisms are isomorphisms of P to its dual poset, that is, bijections σ of P to itself such that $\sigma(x) \leq \sigma(y) \iff x \geq y$.

For any automorphism σ , we say that a lower ideal $E \in J(P)$ is σ -*symmetric* if $\sigma(E) = E$ when σ is increasing, or if $\sigma(E) = P - E$ when σ is decreasing. For any subgroup G of the group of automorphisms, we say that E is G -symmetric if it is σ -symmetric for all $\sigma \in G$. Generally speaking, given P and G we are interested in sampling from the set of G -symmetric lower ideals of P .

When G contains only increasing automorphisms, G -symmetric ideals are in natural bijection with the ideals of the quotient poset P/G , so that one can naturally apply the CFTP algorithm to the quotient poset. The situation is different, however, when G contains decreasing automorphisms. In this case, the quotient set P/G no longer has a naturally defined partial order. Note, however, that quotienting by the subgroup G^+ of increasing automorphisms reduces all increasing automorphisms to the identity, and all decreasing automorphisms to a single decreasing involution. As a result, we can, without loss of generality, assume that the group G only contains the identity and a single decreasing involution σ on P .

Note that if σ has any fixed points, the set of σ -symmetric ideals is empty. In such a case, fixed points are necessarily pairwise incomparable (since σ is decreasing). To obtain the “next best thing” to truly σ -symmetric ideals, we can change the definition of a (quasi) σ -symmetric ideal to $\sigma(E) = P - E - F_\sigma$, where F_σ denotes the set of fixed points of σ . Equivalently, we could replace (P, \leq) by the induced order on $P - F_\sigma$ and sample from the set of “truly” symmetric lower ideals on this smaller set. In the rest of this section, we assume that σ has no fixed points; we will see in Section 3 that this extended definition of symmetric ideals makes sense from a combinatorial point of view when dealing with symmetric ASMs.

2.3 Symmetric CFTP

We now turn to a description of SYMMETRIC-CFTP. In view of the discussion in the previous paragraph, we assume that we are given a finite poset P and a decreasing involution on P , and we are interested in random sampling from the set $J_\sigma(P)$ of (quasi) σ -symmetric lower ideals of P .

A simple observation is the following:

Lemma 1. *Let x and y be two distinct elements of a finite poset P , and let $T_x^\epsilon : J(P) \rightarrow J(P)$ be defined as in Subsection 2.1.*

- if x and y do not cover each other, then T_x^ϵ and T_y^ϵ commute;
- if y does not cover x , then T_y^+ and T_x^- commute.

As a consequence, when x and $\sigma(x)$ are incomparable in the poset P , $S_x^\epsilon = T_x^\epsilon \circ T_{\sigma(x)}^{-\epsilon}$ maps symmetric ideals to symmetric ideals in a similar way that T_x^ϵ maps ideals to ideals; and, as a compose of increasing mappings of $J(P)$ to itself, it is also increasing.

Now if, say, $x > \sigma(x)$, every symmetric ideal contains $\sigma(x)$ but not x . In this situation, S_x^+ is still defined as a commutative compose, and leaves each symmetric ideal unchanged. It may not be true that T_x^- and $T_{\sigma(x)}^+$ commute when x covers $\sigma(x)$, and in this case we define $S_x^- = S_x^+$.

With these definitions, we have, for any two symmetric ideals E and E' , $S_x^\epsilon(E) = E' \iff S_{\sigma(x)}^{-\epsilon}(E') = E$. Thus, provided we can check that the resulting Markov chain on $J_\sigma(P)$ is irreducible (we will see that this is always true), we obtain a CFTP algorithm on $J_\sigma(P)$ by taking the previous MONOTONE-CFTP algorithm on $J(P)$ and simply using “symmetric” update functions S_x^ϵ instead of the original T_x^ϵ .

One possible problem lies in making this algorithm practical. $J_\sigma(P)$ does not have a naturally ordered structure – any symmetric ideal must contain exactly half the elements of P , so no two symmetric ideals can be compared by inclusion. This makes it potentially difficult to detect coalescence of the process.

Fortunately, we still have a naturally ordered space on which the Markov coupling using the S_x^ϵ update functions is monotone, namely, $J(P)$. When run with $J(P)$ as a state space, the Markov chain is no longer irreducible since no transitions lead out of $J_\sigma(P)$ (no non-symmetric state is accessible from any symmetric state), but it is still a monotone grand coupling, and coalescence of the whole state space can still be detected by the condition that the copies started from states \emptyset and P have met.

One can save some time by precomputing a lower and upper bound (in $J(P)$) for $J_\sigma(P)$, and using these two ideals as starting states instead of \emptyset and P . If we define

$$\begin{aligned} E_{\text{low}} &= \{x \in P : x < \sigma(x)\} \\ E_{\text{high}} &= \{x \in P : x \not\prec \sigma(x)\} \end{aligned}$$

then any symmetric ideal contains E_{low} and is included in E_{high} , so that checking coalescence for the coupled chains started at E_{low} and E_{high} is sufficient. Alternately, one can note that $J_\sigma(P)$ is isomorphic to $J_\sigma(E_{\text{high}} - E_{\text{low}})$ (any symmetric ideal of P is the disjoint union of E_{low} and some symmetric ideal of the induced order on $E_{\text{high}} - E_{\text{low}}$) and work on the smaller induced order, where no comparable pairs $(x, \sigma(x))$ remain.

The whole SYMMETRIC-CFTP method can be summed up in the following theorem:

Theorem 1. *Let (P, \leq) be some finite poset, and assume there exists a decreasing involution σ on P with no fixed points. Let P' denote the set of elements of P such that x does not cover, and is not covered by, $\sigma(x)$.*

Then the set $J_\sigma(P)$ of σ -symmetric ideals of P is nonempty. Furthermore, for any probability distribution π on P' such that $\pi(x) = \pi(\sigma(x)) > 0$ for all $x \in P'$, if $(x_k)_{k \geq 1}$ is a sequence of independent random elements of P' with common distribution π , $(\epsilon_k)_{k \geq 1}$ is an independent sequence of independent, uniform random signs, and $S_k = S_{x_1}^{\epsilon_1} \circ \dots \circ S_{x_k}^{\epsilon_k}$, then

- *the random variable $N = \inf\{k : S_k(\emptyset) = S_k(P)\}$ is almost surely finite, with finite expectation;*
- *$S_N(\emptyset)$ is uniformly distributed on $J_\sigma(P)$.*

Proof. We construct a symmetric ideal I explicitly by examining each pair $(x, \sigma(x))$ in an arbitrary order, each time selecting exactly one of the two to be included in I . The rules are as follows:

- if $x < \sigma(x)$, or if I already contains an element y with $y > x$, then add x to I ;

- if $\sigma(x) < x$, or if I already contains an element y with $y > \sigma(x)$, then add $\sigma(x)$ to I ;
- otherwise, arbitrarily select one of x or $\sigma(x)$ to be added to I .

The first and second rules ensure that the set I obtained is indeed a lower ideal of P (which will be σ -symmetric by construction), provided one never runs into a situation where they are in contradiction. Assume that one has $x < y$ and $\sigma(x) < y'$ with both y and y' in I , then one has $\sigma(y') < x < y$ and $\sigma(y) < y'$, so that whichever of y and y' was added last was not added according to the rules; a similar argument holds for situations where x and $\sigma(x)$ are comparable.

The proof that N is almost surely finite and that $S_N(\emptyset)$ does have the required distribution is very similar to the classical proofs for CFTP, and is omitted due to space constraints. □

2.4 A note on coalescence detection

The stopping criterion for the SYMMETRIC-CFTP algorithm, as suggested by Theorem 1, is that the coupled chains started at the minimum and maximum states (or at the lower and higher bounds E_{low} and E_{high} as defined earlier) have coalesced to the same state, which, by monotonicity, implies that the composed update function is constant on the whole of $J(P)$.

This is not necessary, and is not optimal. We only need a stopping criterion which ensures that the composed update function is constant on $J_\sigma(P)$. Since $J_\sigma(P)$ is an antichain in $J(P)$ (no two symmetric ideals are comparable), a sufficient condition for coalescence on $J_\sigma(P)$ is that $S_N(\emptyset) \in J_\sigma(P)$ (or, symmetrically, $S_N(P) \in J_\sigma(P)$). Using this condition in the algorithm will always save time (provided checking symmetry in an ideal is not significantly longer than checking for equality of two ideals), not only because one only has to run one copy of the Markov chain instead of two, but also because this corresponds to running the coupling until the chain started at \emptyset reaches $J_\sigma(P)$, and not until the two chains started at \emptyset and P reach $J_\sigma(P)$.

3 Random sampling from symmetry classes of alternating-sign matrices

In this section, we demonstrate the usefulness of our SYMMETRIC-CFTP framework by applying it to the random generation of alternating-sign matrices with arbitrary symmetry conditions. Though all but one symmetry classes can be tackled using only the classical MONOTONE-CFTP algorithm, our approach has the added benefit of providing a unified treatment of all symmetry classes.

An *alternating-sign matrix* (ASM for short) of size N is an $N \times N$ square matrix whose entries are all 0, 1 or -1 , and such that nonzero entries in each line and column alternate in sign, starting and ending with a 1. These matrices have many fascinating combinatorial properties; see [5] for a survey, or [1] for the story of their enumeration.

There is a simple bijection between ASMs of size N and a set of *height matrices* of size N , which are defined as those $(N + 1) \times (N + 1)$ matrices with integer entries $h_{i,j}$ ($0 \leq i, j \leq N$) satisfying the conditions

- for all $0 \leq i \leq N$, $h_{i,0} = h_{0,i} = h_{N-i,N} = h_{N,N-i} = i$;
- for all $0 \leq i, j < N$, $|h_{i,j} - h_{i,j+1}| = |h_{i,j} - h_{i+1,j}| = 1$.

A simple way of expressing the bijection is as follows: if the entries of an ASM A are $(a_{i,j})_{1 \leq i, j \leq N}$, set

$$h_{i,j} = i + j - 2 \sum_{k \leq i} \sum_{\ell \leq j} a_{i,j}$$

to obtain the corresponding height matrix.

Height matrices of a given size are naturally ordered by entry-wise comparison, and one easily checks that the resulting poset is a distributive lattice where the meet and join operations are defined by taking entry-wise maximums and minimums. The minimum (respectively, maximum) matrices have their entries defined by $h_{i,j} = |i - j|$ (respectively, $h_{i,j} = N - |N - i - j|$). One easily checks that the set of height matrices of size N is isomorphic to the set of ideals of the set

$$P_N = \{(i, j, k) : 1 \leq i, j < N, k = i + j \pmod{2}, |i - j| < k \leq N - |N - i - j|\},$$

with a partial order defined by

$$(i, j, k) \leq (i', j', k') \Leftrightarrow k' \geq k + |i - i'| + |j - j'|.$$

Given a lower ideal E of P_N , the corresponding height matrix has its entries defined by

$$h_{i,j} = \max(|i - j|, \sup\{k : (i, j, k) \in E\}),$$

and the inverse bijection defines the lower ideal corresponding to a height matrix as

$$E = \{(i, j, k) : k \leq h_{i,j}\}.$$

All these bijections make it easy to apply the MONOTONE-CFTP algorithm to the random generation of ASMs. The most basic application of the algorithm would be to select some arbitrary probability distribution on P_N (typically, uniform) and use the update functions T_x^ϵ as we defined them in Section 2; a very simple variant, which is easier to implement in practice, would only ask for a probability distribution on *coordinates* (i, j) (again, the uniform distribution is an obvious and easy choice) and using update functions of the form

$$T_{i,j}^\epsilon = \prod_k T_{(i,j,k)}^\epsilon.$$

(In the above formula, the product denotes composition; note that all involved $T_{i,j,k}^\epsilon$ update functions commute, so that the notation is not ambiguous.)

In their matrix formulation, the update functions $T_{i,j}^+$ (respectively, $T_{i,j}^-$) correspond to the instructions “increase (respectively, decrease) entry $h_{i,j}$ by 2 if this is compatible with the height matrix conditions, otherwise do nothing”, and seem more natural than the $T_{i,j,k}^\epsilon$, which correspond to “increase (decrease) entry $h_{i,j}$ to k if possible, otherwise do nothing”.

The 8-element dihedral group of symmetries of the square acts naturally on ASMs by permuting entries, so that, for each of its subgroups, one may consider the set of ASMs that are invariant under its action. It turns out that each of these symmetries corresponds, through the standard isomorphism we described above, to an (increasing or decreasing) automorphism of the poset P_N , so that either classical MONOTONE-CFTP or SYMMETRIC-CFTP can be used to sample from each symmetry class. In some cases, no truly symmetric ASMs exist for some sizes, due to the existence of fixed points for the decreasing automorphisms; in such cases, we instead turn to quasi-symmetric ASMs.

3.1 Diagonal symmetry

Height matrices can be symmetric around their main diagonal ($h_{i,j} = h_{j,i}$ for all i, j), which corresponds exactly to symmetric ideals for the (increasing) automorphism

$$\sigma_D : (i, j, k) \mapsto (j, i, k).$$

Symmetry around the other diagonal corresponds to the automorphism

$$\sigma_A : (i, j, k) \mapsto (N - j, N - i, k).$$

Diagonally and doubly-diagonally symmetric ASMs of all sizes exist, though no formula for their enumeration is known (and no simple formula is likely to exist, since the enumerating sequence does not seem to factor into small primes).

3.2 Vertical or horizontal symmetry

ASMs invariant under symmetry around a vertical axis correspond to height matrices satisfying $h_{i,N-j} = N - h_{i,j}$ for all i, j . The corresponding involution on P_N is

$$\sigma_V : (i, j, k) \mapsto (i, N - j, N + 2 - k)$$

which is a *decreasing* automorphism. Note that, when N is even, σ_V has fixed points of the form $(i, N/2, N/2 + 1)$ with odd i , so that truly vertically symmetric ASMs of even sizes do not exist (this is, of course, very easy to see directly in the ASM formulation: since every ASM has a single nonzero entry in its first line, this entry must be in the middle column for vertical symmetry, and even sized matrices do not have a middle column).

If we turn to quasi-vertically symmetric ASMs of even sizes, these correspond to σ_V -symmetric ideals of P_N with the fixed points removed. The corresponding height matrices have entries in their middle column alternate between $N/2$ and $N/2 - 1$. These were counted by Kuperberg [4] under the guise of UASMs (U-turn ASMs), with a slight change of size; there is an easy bijection between quasi-vertically symmetric ASMs of size $2N$ and Kuperberg's UASMs of size $2N - 2$.

Horizontal symmetry is of course similar, corresponding to the decreasing automorphism

$$\sigma_H : (i, j, k) \mapsto (N - i, j, N - k).$$

For the random generation of vertically symmetric ASMs, SYMMETRIC-CFTP can be used. Note, however, that the set of fixed entries in height matrices (the two middle columns for odd size; the middle column and one entry in every two in adjoining columns for even sizes) separates nonfixed entries in the left and right halves of the matrix. This makes it possible to define a distributive lattice structure on the set of (quasi-)symmetric ASMs by using only the left half of the height matrix columns for entry-wise comparison; with this *ad hoc* partial ordering, classical MONOTONE-CFTP can also be used (and is actually the same algorithm) in this case.

ASMs that are both horizontally and vertically symmetric can be treated in a similar way, using both σ_H and σ_V . Again, one can use MONOTONE-CFTP by using only the top left quarter of the height matrices for comparisons. Generating functions for quasi-symmetric ASMs of even size $4n$ were described by Kuperberg under the name of UUASMs (double U-turn ASMs) of size $4n - 4$. Quasi-symmetric ASMs of size $4n + 2$ seem to be missing from current literature; their enumeration sequence, starting with 1, 2, 28, 3146, 2855320, does not appear in [10].

3.3 Rotational symmetry

Half-turn symmetric ASMs were enumerated by Kuperberg [4] for even size, and by Razumov and Stroganov [7] for odd size. The symmetry condition on height matrices ($h_{N-i, N-j} = h_{i, j}$) corresponds to σ_C -symmetry for the increasing automorphism

$$\sigma_C : (i, j, k) \mapsto (N - i, N - j, k),$$

and again classical MONOTONE-CFTP can be used to sample from this symmetry class.

Quarter-turn symmetry corresponds to the condition $h_{j, N-i} = N - h_{i, j}$ for height matrices, or to σ_Q -symmetry for the decreasing automorphism

$$\sigma_Q : (i, j, k) \mapsto (j, N - i, N + 2 - k).$$

This automorphism has a single fixed point $(N/2, N/2, N/2 + 1)$ if N is of the form $4n + 2$, and none in other cases. Truly quarter-turn-symmetric ASMs of size divisible by 4 were enumerated by Kuperberg [4], and those with odd size by Razumov and Stroganov [8]; a similar enumeration formula for quasi-quarter-turn-symmetric ASMs with size $4n + 2$ is conjectured in [3].

For all sizes of quarter-turn (quasi-)symmetric ASMs, classical MONOTONE-CFTP does not seem to be adapted to random sampling, as there is no natural way to define a partial ordering on such ASMs. In this case, SYMMETRIC-CFTP is a natural and efficient choice.

3.4 Full symmetry

The maximum amount of symmetry one can ask for in ASMs is invariance under the action of the whole group of symmetries of the unit square. In the P_N formulation, this corresponds to (quasi-)symmetry under σ_D , σ_V , and their composites. As was the case for vertical symmetry, despite the fact that a decreasing automorphism is used, the whole symmetry class can be turned into a distributive lattice by using only a triangular-shaped region of the matrix (*e.g.*, entries $h_{i, j}$ with $i \leq j \leq N/2$) to define the partial ordering, so that an ad hoc version of MONOTONE-CFTP can be used instead of SYMMETRIC-CFTP.

3.5 Simulation results

We implemented our algorithm to generate random ASMs with prescribed symmetry conditions, and some results with moderate size are presented in Figure 1. White dots represent positive entries, black dots represent negative entries, and gray areas cover zones where the symmetry constraints are not respected due to quasi-symmetry; one sample from each symmetry class is shown. For large sizes, one observes that nonzero entries are very rare outside of a roughly circular shape, a phenomenon known as the *arctic circle phenomenon*; the exact shape (or even existence) of this frozen region has not yet been proved for uniform random ASMs, though Colomo and Pronko [2] recently announced a partial proof that this limit shape exists and is not circular.

We also give experimental results on coalescence times (in number of time steps) in Table 1. We ran the coupled copies of the chain, starting from the minimum and maximum states (or from E_{low} and E_{high} when appropriate), forward in time until both met. While the resulting states would not be uniform, the distribution of the coalescence time is the same as when the chain is run backward, and this takes less simulation time. This number of steps is a good measure for the running time of the CFTP algorithm, which will have to perform a number of simulation steps roughly proportional to it, both in the “binary-backoff” and “read-once” [12] variants.

All simulations were done using symmetric moves of the form $S_{i, j}^\epsilon = \prod_k S_{i, j, k}^\epsilon$ for the definition of $S_{i, j, k}^\epsilon$ appropriate to the considered symmetry class, using the corresponding

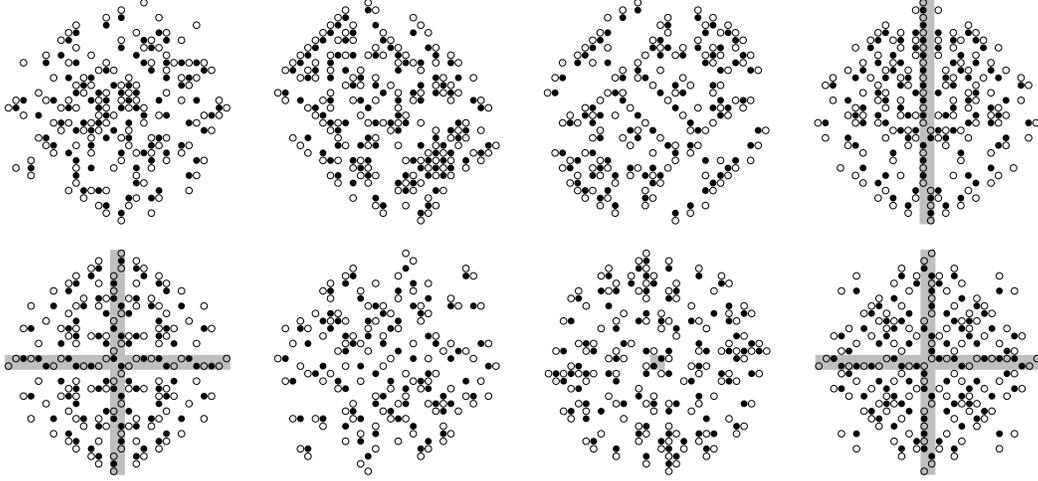


Figure 1: Random ASMs of size 30

E_{low} and E_{high} as starting points. The coordinates (i, j) were taken uniformly at random inside the height matrix. The averages and standard deviations are over 300 runs for each class for size 10, 100 runs for sizes 20 and 40, and 20 runs for size 60. The average coalescence times seem to be close to $C.N^\alpha$, with $\alpha \simeq 4.4$ and C depending on the symmetry class.

	$N = 10$	$N = 20$	$N = 40$	$N = 60$
No symmetry	11908±4197	289741±69692	$5.9 \cdot 10^6 \pm 1.1 \cdot 10^6$	$3.4 \cdot 10^7 \pm 8.1 \cdot 10^6$
Diagonal	6528±2808	144917±51115	$2.8 \cdot 10^6 \pm 7.1 \cdot 10^5$	$1.6 \cdot 10^7 \pm 4.0 \cdot 10^6$
Double diag.	3481±1561	71557±25490	$1.4 \cdot 10^6 \pm 3.6 \cdot 10^5$	$7.6 \cdot 10^6 \pm 1.2 \cdot 10^6$
Vertical	1201±582	34881±12312	$8.2 \cdot 10^5 \pm 2.2 \cdot 10^5$	$5.2 \cdot 10^6 \pm 1.4 \cdot 10^6$
Vert. & horiz.	159±113	6507±2815	$1.7 \cdot 10^5 \pm 5.8 \cdot 10^4$	$1.0 \cdot 10^6 \pm 2.7 \cdot 10^5$
Half turn	5546±2403	127685±37860	$2.7 \cdot 10^6 \pm 6.5 \cdot 10^5$	$1.5 \cdot 10^7 \pm 3.4 \cdot 10^6$
Quarter turn	1054±545	28938±11567	$7.1 \cdot 10^5 \pm 2.0 \cdot 10^5$	$4.3 \cdot 10^6 \pm 1.1 \cdot 10^6$
Total symmetry	91±69	3385±1810	$8.1 \cdot 10^4 \pm 2.8 \cdot 10^4$	$4.7 \cdot 10^5 \pm 1.5 \cdot 10^5$

Table 1: Observed coalescence times (average and standard deviation)

In cases where the symmetry class is defined through a group with only increasing automorphisms, the coalescence time seems to be roughly divided by the order of the group between nonsymmetric and symmetric ASMs of the same size. This relationship also holds between the vertically and horizontally symmetric and totally symmetric cases. In cases where the group contains negative automorphisms, no such simple relationship seems to hold; the situation is made more complex by the change in the starting states.

4 Concluding remarks

In this paper, we have described a new variant of the MONOTONE-CFTP algorithm, which is well suited to sampling from some symmetric subsets of finite distributed lattices. While we have mainly studied its application to generating random symmetric alternating-sign matrices, it can be applied to other similar situations; an example which immediately

comes to mind is that of symmetric plane partitions (the fact that the enumeration sequences for symmetry classes of ASMs and plane partitions often derive from each other is irrelevant here).

Strictly speaking, (quarter-turn) rotational invariance is the only symmetry class for which one cannot adapt the usual MONOTONE-CFTP algorithm, but even for other symmetry classes, our setting has the advantage of giving a unified treatment.

On the theoretical side, an unsolved issue is that of the coalescence time. In all the examples we have experimented with, the coalescence time for sampling from the symmetric ideals is significantly lower than that for sampling from the whole ambient lattice. It would be interesting to give a precise and rigorous statement of this observation, and to know whether this is a general phenomenon or is due to the special structure of the posets we studied.

Acknowledgements

This research was supported by French ANR project SADA (ANR-05-BLAN-0372).

References

- [1] David Bressoud. *Proofs and confirmations: the story of the alternating-sign matrix conjecture*. Cambridge University Press, 1999.
- [2] Filippo Colomo and Andrey Georgievich Pronko. Email communication on the domino forum. See also [arXiv:math-ph/0704.0362](https://arxiv.org/abs/math-ph/0704.0362), 2008.
- [3] Philippe Duchon. On the link pattern distribution of quarter-turn symmetric FPL configurations. To appear in FPSAC 2008. Available as [arXiv:math.CO/0711.2871](https://arxiv.org/abs/math.CO/0711.2871).
- [4] Greg Kuperberg. Symmetry classes of alternating-sign matrices under one roof. *Annals of Mathematics*, 156:835–866, 2002.
- [5] James Propp. The many faces of alternating-sign matrices. In *Discrete Models: Combinatorics, Computation, and Geometry*, volume AA of *DMTCS Proceedings*, pages 43–58, 2001.
- [6] James Propp and David Wilson. Exact sampling with coupled Markov chains and applications to statistical physics. *Random Structures and Algorithms*, 9:223–252, 1996.
- [7] Alexander Razumov and Yuri Stroganov. Combinatorial nature of ground state of $O(1)$ loop model. *Theor. Math. Phys.*, 138:333–337, 2004.
- [8] Alexander Razumov and Yuri Stroganov. Enumeration of quarter-turn-symmetric alternating sign matrices of odd order. *Theoretical and Mathematical Physics*, 149(3):1639–1650, 2006.
- [9] David Robbins. Symmetry classes of alternating sign matrices. Unpublished manuscript [arXiv:math.CO/000804](https://arxiv.org/abs/math.CO/000804).
- [10] Neil Sloane. The on-line encyclopedia of integer sequences. Available online at <http://www.research.att.com/~njas/sequences/>.
- [11] Richard Stanley. *Enumerative Combinatorics, volume 2*. Number 62 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999.
- [12] David Wilson. How to couple from the past using a read-once source of randomness. *Random Structures and Algorithms*, 16(1):85–113, 2000.