

COUPLING ANY NUMBER OF BALLS IN THE INFINITE-BIN MODEL

KSENIA CHERNYSH,* *Heriot-Watt University*

SANJAY RAMASSAMY,** *Brown University*

Abstract

The infinite-bin model, introduced by Foss and Konstantopoulos (2003), describes the Markovian evolution of configurations of balls placed inside bins, obeying certain transition rules. We prove that we can couple the behaviour of any finite number of balls, provided at least two different transition rules are allowed. This coupling makes it possible to define the regeneration events needed by Foss and Zachary (2013) to prove convergence results for the distribution of the balls.

Keywords: Markov chain; coupling; infinite-bin model

2010 Mathematics Subject Classification: Primary 60J10
Secondary 68R05

1. Introduction

1.1. Background

Interacting particle systems and their stochastic evolution have been widely studied in probability and statistical physics. One such system, the infinite-bin model, was introduced by Foss and Konstantopoulos [2] as an abstraction of stochastic ordered graphs, which are a directed version of Erdős–Rényi graphs that have applications in queuing theory, mathematical ecology, and performance evaluation of computer systems (see [2] and the references therein). A very similar (though less general) model had also appeared in earlier work of Aldous and Pitman [1].

In the infinite-bin model, a configuration is made up of infinitely many bins indexed by the nonpositive integers, each bin containing a positive and finite number of balls. An elementary random move consists in picking one ball at random (according to a certain probability distribution) and adding one ball to the bin situated immediately to the right of the ball we picked. If the ball we picked was already in the rightmost bin, we create a new bin to its right, add a ball in it, and relabel the bins so that the new rightmost bin has label 0. The stochastic dynamics arises from the iteration of independent and identically distributed (i.i.d.) elementary random moves.

Questions of interest include the existence and uniqueness of a stationary solution, the convergence to the stationary solution for an arbitrary initial configuration, and the rate of creation of new bins. The first two questions have been tackled in certain cases in [2] and [3]. The last one was discussed in some cases in [2] and [4].

Received 8 July 2016; revision received 10 December 2016.

* Postal address: School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK.

** Current address: Unité de Mathématiques Pures et Appliquées, École normale supérieure de Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France. Email address: ramassam@math.brown.edu

In this paper we prove that in all the nontrivial cases, we can find a sequence of moves such that, after applying that sequence of moves, the position of the rightmost N balls is some prescribed configuration, independent of the initial configuration (here N is an arbitrary positive integer).

This is a result needed by Foss and Zachary [3] to prove the convergence to the stationary distribution. They do not prove it in their paper, but instead point to the present paper for a proof. Note that the present paper relaxes the assumptions of [3] so that we can now cover all the cases where the theory of renovation events works.

We finally mention that we can rephrase this result in the framework of automata theory. The infinite-bin model (or rather the finite projections of it defined in the next subsection) can be seen as a finite automaton, and the sequence of moves we construct corresponds to a synchronizing word in the language of automata theory (see, e.g. [5]).

1.2. Definitions and main result

An *infinite configuration* X is a sequence of positive integers $(X(i))_{i \in \mathbb{Z}_-}$ indexed by $\mathbb{Z}_- = \{0, -1, -2, \dots\}$. The integer $X(i)$ represents the number of balls in the bin labeled by i . We adopt the unusual indexing by nonpositive integers to conform to the original definition of the infinite-bin model in [2]. It finds its roots in an application to stochastic ordered graphs.

For any integer $k \geq 1$, we are going to construct the *move of type k* as a function ϕ_k from the set of infinite configurations to itself. Fix an infinite configuration X . The infinite configuration $Y = \phi_k(X)$ is defined as follows. If $X(0) \geq k$, we set

$$Y(i) = \begin{cases} 1 & \text{if } i = 0, \\ X(i + 1) & \text{if } i < 0. \end{cases}$$

If $X(0) < k$, we define

$$i_k = \min \left\{ j \in \mathbb{Z}_- \mid \sum_{i=j}^0 X(i) < k \right\}$$

and set

$$Y(i) = \begin{cases} X(i) + 1 & \text{if } i = i_k, \\ X(i) & \text{if } i \neq i_k. \end{cases}$$

In words, the move of type k adds one ball inside the bin situated immediately to the right of the bin containing the k th ball, where the balls are counted from right to left. When the k th ball is already in the rightmost bin, we create a new bin immediately to its right, add a ball in it, and relabel all the bins in such a way that the newly created bin will be labeled by 0. See Figure 1 for an example.

It will be useful to consider configurations with finitely many balls. If n is a positive integer, we define an n -*configuration* X to be a sequence of nonnegative integers $(X(i))_{i \in \mathbb{Z}_-}$ such that

- $\sum_{i \in \mathbb{Z}_-} X(i) = n$;
- there exists an integer $p \geq 1$ such that $X(i) > 0$ if $i > -p$ and $X(i) = 0$ otherwise.

We will then denote the n -configuration by the p -tuple $[X(-p + 1), \dots, X(0)]$, omitting the infinite string of 0s on the left.

If X is an infinite configuration and n is a positive integer, we define $\pi_n(X)$ the n -*ball projection* of X to be the n -configuration obtained from X by keeping the rightmost n balls and

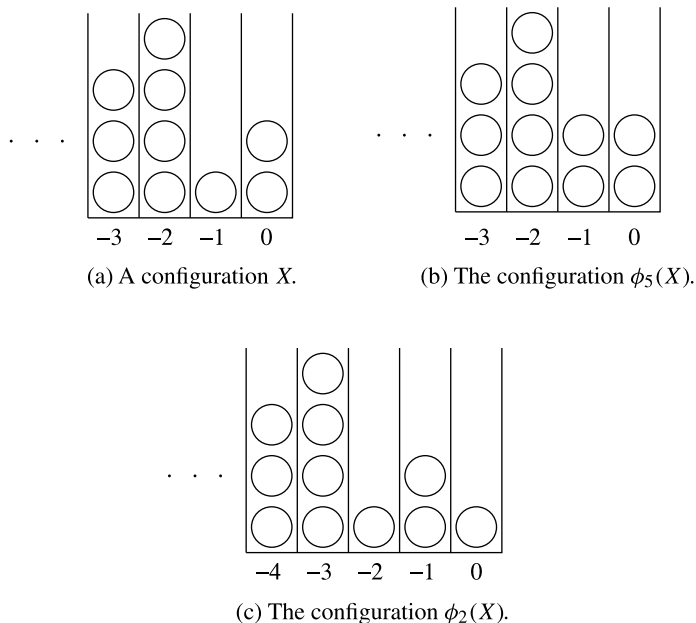


FIGURE 1.

erasing all the other balls. More precisely,

$$(\pi_n(X))_i = \begin{cases} \min(X(i), n - \sum_{j=i+1}^0 X(j)) & \text{if } \sum_{j=i+1}^0 X(j) < n, \\ 0 & \text{otherwise.} \end{cases}$$

An *algorithm* Φ is defined as the composition of a finite number of moves. It maps one infinite configuration to another. The *length* of an algorithm is the number of moves used to build it. For example, $\Phi = \phi_5^3\phi_2$ is an algorithm of length 4 obtained by applying ϕ_2 first followed by three times ϕ_5 .

An algorithm Φ is said to *couple the first n balls* if, for any two infinite configurations X and Y , we have $\pi_n(\Phi(X)) = \pi_n(\Phi(Y))$. We have the following result regarding the existence of coupling algorithms.

Theorem 1. *For any integers $N \geq 1$ and $1 \leq k < l$, we can find an algorithm using only the moves ϕ_k and ϕ_l which couples the first N balls. The length of the algorithm can be taken to be less than $N + 4l^2$.*

The proof is provided in Section 2. We will see in the proof that we can impose the condition that any infinite configuration obtained after applying that coupling algorithm contains at least k balls in the rightmost bin, which is exactly the setting in which Lemma 2 of [3] was stated. Note that we do not need the condition (required by that Lemma 2) that k and l are co-prime.

1.3. Probabilistic implications

Coupling algorithms are interesting from a probabilistic point of view because they make it possible to define regeneration events, and thus make the infinite-bin model more tractable.

Let ξ be a random variable taking values in the set of the positive integers, and let $(\xi_n)_{n \geq 0}$ be a sequence of i.i.d. random variables distributed like ξ . Then any sequence of infinite configurations $(X_n)_{n \geq 0}$ such that, for all $n \geq 0$, $X_{n+1} = \phi_{\xi_n}(X_n)$ is a Markov chain.

Theorem 4 in [3] proves, under some assumptions, the convergence of the distribution of the number of balls in the rightmost k bins for any positive integer k . It is based on their Lemma 2, which is a more restrictive version of our Theorem 1. For a proof of that lemma, they point to the present paper. Combining our Theorem 1 with Theorem 4 and Proposition 2 of [3], leads us to the following corollary.

Corollary 1. (See Foss and Zachary [3].) *Assume that ξ has finite expectation and is not constant almost surely (a.s.). Then for any integer $k \geq 0$ and any initial infinite configuration $X_0, (X_n(-k), \dots, X_n(0))$ converges to a proper limiting random vector in the total variation norm. Therefore, X_n weakly converges to its proper limit.*

Here ‘proper’ means the limiting random variable is finite a.s. Note that if ξ is a.s. constant equal to c then the dynamics is deterministic and ultimately c -periodic.

The proof is the same as that of Theorem 4 in [3], replacing their Lemma 2 by our Theorem 1 and constructing regeneration events. A regeneration event E is an event such that at the time when E starts, conditionally on E , the future does not depend on what happened before E (here ‘depend’ is used in the algebraic sense, not in the probabilistic sense). They define a regeneration event to be a coupling algorithm followed by an infinite sequence of moves with the property that the balls involved in that infinite sequence of moves are only those balls that were created by the coupling algorithm. Such an event will eventually occur a.s. This implies that if we run the Markov chain with two different initial configurations but with the same sequence $(\xi_n)_{n \geq 0}$, the content of the first k bins of the Markov chains will eventually be the same a.s.

An interesting next step would be to find the coupling algorithms that occur most frequently, i.e. for which the ratio ‘probability of algorithm’ divided by ‘length of algorithm’ is maximal. When the distribution of ξ is the uniform measure on a finite set of integers, this boils down to finding the shortest coupling algorithms, which is a classical question in automata theory (see, e.g. [5]).

2. Proof of Theorem 1

The structure of the proof is as follows. Firstly, we reduce the proof to finding an algorithm which couples the first k balls when applied to l -configurations. For this purpose, we will need to define how moves act on l -configurations. Secondly, we define a set of k l -configurations $\{X_0, \dots, X_{k-1}\}$ and two algorithms Ψ_1 and Ψ_2 using only the moves ϕ_k and ϕ_l verifying the following properties:

- for any l -configuration X , $\Psi_1(X) \in \{X_0, \dots, X_{k-1}\}$;
- for any $0 \leq i \leq k - 1$, $\Psi_2(X_i) = X_0$.

It will follow that the algorithm $\Psi_2\Psi_1$ couples the first l balls (and *a fortiori* the first k balls) when applied to l -configurations.

The $k = 1$ case is easy: applying move ϕ_1 N times couples the first N balls. From now on we will assume that $k \geq 2$.

If $n \geq k$ and X, Y are two infinite configurations such that $\pi_n(X) = \pi_n(Y)$, observe that $\pi_{n+1}(\phi_k(X)) = \pi_{n+1}(\phi_k(Y))$. So without loss of generality, we may assume that $N = k$: if

we find an algorithm coupling the first k balls, applying $N - k$ more times the move ϕ_k will couple the first N balls.

We now define projections and moves applied to n -configurations. If $n > m$ one defines $\pi_m(X)$, the m -ball projection of an n -configuration X in the same way as for infinite configurations. The map π_m is a function from n -configurations into m -configurations. If X is an n -configuration and $k \leq n$, we define $\phi_k(X)$ to be the following n -configuration: starting from X , add one ball to the bin situated immediately to the right of the bin containing the k th ball (where the balls are counted from right to left), then delete one ball from the leftmost bin. When the k th ball is already in the rightmost bin, we create a new bin immediately to its right, add a ball in it and relabel all the bins in such a way that the newly created bin will be labeled by 0. The map ϕ_k is a function from n -configurations into n -configurations. This definition of ϕ_k for n -configurations parallels the one for infinite configurations, so that moves commute with projections, i.e. if $k \leq n$ and if X is an infinite configuration or an m -configuration with $m \geq n$, then

$$\pi_n(\phi_k(X)) = \phi_k(\pi_n(X))$$

This is illustrated in Figure 2.

Fix an integer $n \geq 1$. If, for any move ϕ_k entering in the construction of an algorithm Φ , we have $k \leq n$ then Φ also maps n -configurations to n -configurations.

Let Φ be an algorithm using only the moves ϕ_k and ϕ_l which couples the first k balls when restricted to l -configurations. Then, if X and Y are two infinite configurations,

$$\begin{aligned} \pi_k(\Phi(X)) &= \pi_k(\pi_l(\Phi(X))) \\ &= \pi_k(\Phi(\pi_l(X))) \\ &= \pi_k(\Phi(\pi_l(Y))) \\ &= \pi_k(\pi_l(\Phi(Y))) \\ &= \pi_k(\Phi(Y)), \end{aligned}$$

where the third equality follows from the fact that Φ couples the first k balls of the l -configurations $\pi_l(X)$ and $\pi_l(Y)$.




FIGURE 2: Commutation of ϕ_5 with π_8 .

So Φ also couples the first k balls for infinite configurations. Thus, we can reduce it to the question of finding an algorithm Φ which couples the first k balls when applied to l -configurations. From now on, by ‘configurations’ we will mean ‘ l -configurations’.

Write $l = kd + r$, where d and r are two integers such that $d \geq 1$ and $1 \leq r \leq k$. Define the following configurations. Set $X_0 = [k, \dots, k, r]$ with d consecutive bins containing k balls and the rightmost bin containing r balls. If $1 \leq i \leq r - 1$, set $X_i = [i, k, \dots, k, r - i]$, with d consecutive bins containing k balls sandwiched between a bin with i balls and a bin with $r - i$ balls. If $r \leq i \leq k - 1$, set $X_i = [i, k, \dots, k, k + r - i]$, with $d - 1$ consecutive bins containing k balls sandwiched between a bin with i balls and a bin with $r - i$ balls. Then, for any $1 \leq i \leq k$, $\phi_k(X_i) = X_{i-1}$, where indices are taken modulo k .

Example 1. In the case when $k = 3$ and $l = 11$, we have $d = 3$ and $r = 2$. Thus, $X_0 = [3, 3, 3, 2]$, $X_1 = [1, 3, 3, 3, 1]$, and $X_2 = [2, 3, 3, 3]$.

The rest of the proof of the theorem is based on three lemmas, which will be proved in Section 3. Set

$$M = \max\left(l, \frac{k(k-1)}{2}\right)$$

and $\Psi_1 = \phi_k^M$, where ϕ_k^M denotes the algorithm built from the move k composed M times. Then we have the following result.

Lemma 1. For any l -configuration X , $\Psi_1(X) \in \{X_0, \dots, X_{k-1}\}$.

Define the algorithm

$$\Psi = \phi_k^{l-k} \phi_l^{dr+k[d(d-1)/2]} \phi_k^{k-r}.$$

If $n \in \mathbb{Z}$ is an integer, we write $(n)_+ = \max(n, 0)$. Define the function f from the set $\{0, \dots, k - 1\}$ to itself by

$$f(i) = \begin{cases} (k - l + di)_+ & \text{if } 0 \leq i \leq k - r - 1, \\ (k - (d + 1)(k - i))_+ & \text{if } k - r \leq i \leq k - 1. \end{cases}$$

Using the function f , we can describe the action of Ψ on the configurations X_i as follows.

Lemma 2. If $0 \leq i \leq k - 1$ then $\Psi(X_i) = X_{f(i)}$.

Lemma 3. The function f has the following properties:

- $f(0) = 0$;
- for any $1 \leq i \leq k - 1$, we have $f(i) < i$.

From Lemma 3, we deduce that for any $0 \leq i \leq k - 1$, we have $f^{k-1}(i) = 0$. So by Lemma 2, for any $0 \leq i \leq k$, we have $\Psi^{k-1}(X_i) = X_0$. Combining this with Lemma 1, we conclude that for any l -configuration X , we have $\Psi_2 \Psi_1(X) = X_0$, provided we define $\Psi_2 = \Psi^{k-1}$.

The configuration X_0 has k balls in the rightmost bin, confirming the claim made after the statement of Theorem 1 in Section 1. Also, note that we did better than coupling the first k balls, we coupled the first l balls.

Example 2. Continuing with Example 1, when $k = 3$ and $l = 11$, we have $M = 11$ and $\Psi = \phi_3^8 \phi_{11}^{15} \phi_3$, thus, $\Psi_1 = \phi_3^{11}$ and $\Psi_2 = \Psi^2 = \phi_3^8 \phi_{11}^{15} \phi_3^9 \phi_{11}^{15} \phi_3$. Consider the 11-configuration

$X = [2, 3, 1, 3, 2]$. We first have $\Psi_1(X) = [1, 3, 3, 3, 1] = X_1$, then $\Psi(X_1) = [3, 3, 3, 2] = X_0$ and $\Psi(X_0) = X_0$, thus, $\Psi_2\Psi_1(X) = X_0$.

The length of $\Psi_2\Psi_1$ is

$$L = \max\left(l, \frac{k(k-1)}{2}\right) + (k-1)\left(k-r+dr+k\frac{d(d+1)}{2}+l-k\right).$$

Using the fact that $r \leq k \leq l$ and $d \leq l/k$, we obtain

$$\begin{aligned} L &\leq l + \frac{k(k-1)}{2} + k\left(dr + dk + k\frac{d(d-1)}{2} + l\right) \\ &\leq l + \frac{k^2}{2} + k\left(l + l + k\frac{l^2}{2k^2} + l\right) \\ &\leq l + \frac{l^2}{2} + 3kl + \frac{l^2}{2} \\ &\leq l + 4l^2. \end{aligned}$$

If we want to couple N balls with $N > l$ for infinite configurations, we will need at most $N - l$ more iterations of ϕ_l following Φ . Hence, an upper bound of the total length of such a coupling algorithm is $N + 4l^2$. This concludes the proof of the theorem. \square

3. Proof of the lemmas

Proof of Lemma 1. We will use the following characterization of the configurations X_i : if an l -configuration X has at most k balls in each bin, and every bin except maybe for the leftmost and the rightmost bins contains exactly k balls, then $X \in \{X_0, \dots, X_{k-1}\}$. Also, in the proof of this lemma, it might be convenient to think that the bins do not get relabeled after a new rightmost bin gets created by a move.

Fix an l -configuration X and write $\pi_k(X) = [i_1, i_2, \dots, i_p]$ for some positive integers i_1, \dots, i_p . After applying i_1 times ϕ_k to X , the bin of X that initially contained i_2 balls now contains $i_1 + i_2$ balls. Iterating this process, we observe that if we set

$$n = \sum_{j=1}^{p-1} (p-j)i_j,$$

then after applying n times ϕ_k to X , the rightmost bin of X (which initially contained i_p balls) contains $i_1 + \dots + i_p = k$ balls. All further applications of ϕ_k will add new bins to the right of the original rightmost bin of X . These new bins will all contain k balls, except maybe for the last (rightmost) bin, which will contain at most k balls.

Note that we can write

$$n = \sum_{i=1}^k (\text{distance of the bin containing ball number } i \text{ to the rightmost bin}),$$

where the balls are counted from right to left. The distance between two bins is given by the absolute value of the difference of their labels. That sum is clearly maximal when $p = k$ and $i_1 = \dots = i_k = 1$, thus, $n \leq \frac{1}{2}k(k-1)$.

Recall that we had

$$M = \max(l, \frac{1}{2}k(k - 1)).$$

Since $M \geq l$, all the balls that were originally in X have been deleted by applying ϕ_k^M . Since $M \geq n$, after applying ϕ_k^M , all the bins to the right of the original rightmost bin will contain k balls, except maybe for the leftmost and the rightmost, which will nevertheless contain at most k balls. So $\phi_k^M(X) \in \{X_0, \dots, X_{k-1}\}$. \square

Proof of Lemma 2. If $1 \leq j \leq l - 1$, set $Y_j = [j, l - j]$ and set $Y_0 = [l]$. For any $1 \leq j \leq l$, $\phi_l(Y_j) = Y_{j-1}$, where indices are taken modulo l . We will be using the following lemma, describing how powers of ϕ_k and ϕ_l act on certain configurations.

Lemma 4. (i) Assume that i_1, \dots, i_p are p positive integers summing to l and $n = \sum_{j=1}^{p-2} (p - 1 - j)i_j$. Then

$$\phi_l^n([i_1, i_2, \dots, i_p]) = Y_{l-i_p}.$$

(ii) For any $0 \leq j \leq l - 1$, we have

$$\phi_k^{l-k}(Y_j) = X_{(k-l+j)_+}.$$

Proof. The first part is obtained by iterating the following procedure to remove the leftmost bin:

$$\phi_l^{i_1}([i_1, i_2, \dots, i_p]) = [i_1 + i_2, i_3, \dots, i_p].$$

For the second part, $Y_j = [j, l - j]$, with the leftmost bin being empty if $j = 0$. We need to distinguish three cases.

- Assume that $k - l + j \leq 0$. Then the rightmost bin already contains at least k balls, so $\phi_k^{l-k}([j, l - j]) = X_0 = X_{(k-l+j)_+}$.
- Assume that $0 < k - l + j \leq l - k$. Then $k - l + j$ and $2l - 2k - j$ are both nonnegative. We first compute

$$\phi_k^{k-l+j}([j, l - j]) = [l - k, k].$$

Since $2l - 2k - j \geq l - 2k$, applying $\phi_k^{2l-2k-j}$ to $[l - k, k]$ will delete at least $l - 2k$ balls from the leftmost bin (initially containing $l - k$ balls) and will construct columns of size k , except the rightmost column which will contain at most k balls. So $\phi_k^{2l-2k-j}([l - k, k])$ has to be equal to some X_i . The number of balls in the rightmost bin of X_i will be congruent to $2l - 2k - j = (2d - 2)k + 2r - j$ modulo k , so the number of balls in its leftmost bin will be congruent to $j - r = k - l + j + (d - 1)k$ modulo k . So $\phi_k^{2l-2k-j}([l - k, k]) = X_{k-l+j}$. So

$$\phi_k^{l-k}([j, l - j]) = X_{k-l+j} = X_{(k-l+j)_+}.$$

- Assume that $k - l + j > l - k$. Then

$$\phi_k^{l-k}([j, l - j]) = [j + k - l, 2l - k - j],$$

because the rightmost bin will always have less than k balls in the process of applying $l - k$ times ϕ_k . Moreover, the leftmost bin of $[j + k - l, 2l - k - j]$ also has less than k balls, so that configuration has to be one of the X_i , namely X_{k-l+j} . Thus, $\phi_k^{l-k}([j, l - j]) = X_{(k-l+j)_+}$.

This completes the proof. \square

Let us now use it to prove Lemma 2. We need to distinguish two cases.

- Assume that $0 \leq i \leq k - r - 1$. Then $\phi_k^{k-r}(X_i) = X_{i+r}$, with $r \leq i + r \leq k - 1$. So $X_{i+r} = [i + r, k, \dots, k, k - i]$, with $d - 1$ bins containing k balls sandwiched between a bin with $i + r$ balls and a bin with $k - i$ balls. If we set

$$m_1 = (i + r)(d - 1) + \sum_{j=1}^{d-2} kj,$$

then by the first part of Lemma 4 we obtain

$$\phi_l^{m_1}([i + r, k, \dots, k, k - i]) = Y_{l-k+i}.$$

Set

$$m_2 = dr + k \frac{d(d - 1)}{2} - m_1 = l - k + i - di.$$

Then we have

$$\phi_l^{m_2}(Y_{l-k+i}) = Y_{di}.$$

By the second part of Lemma 4, we obtain

$$\phi_k^{l-k}(Y_{di}) = X_{(k-l+di)_+}.$$

Putting everything together, we conclude that $\Psi(X_i) = X_{f(i)}$ in this case.

- Assume that $k - r \leq i \leq k - 1$. Then $\phi_k^{k-r}(X_i) = X_{i+r-k}$ with $0 \leq i + r - k \leq r - 1$. So $X_{i+r-k} = [i + r - k, k, \dots, k, k - i]$ with d bins containing k balls sandwiched between a bin with $i + r - k$ balls and a bin with $k - i$ balls. If we set

$$n_1 = d(i + r - k) + \sum_{j=1}^{d-1} kj,$$

then by the first part of Lemma 4, we obtain

$$\phi_l^{n_1}([i + r - k, k, \dots, k, k - i]) = Y_{l-k+i}.$$

Set

$$n_2 = dr + k \frac{d(d - 1)}{2} - n_1 = d(k - i).$$

Then we have

$$\phi_l^{n_2}(Y_{l-k+i}) = Y_{r-k+i(d+1)}.$$

By the second part of Lemma 4, we obtain

$$\phi_k^{l-k}(Y_{r-k+i(d+1)}) = X_{(k-l+r-k+i(d+1))_+} = X_{(k-(d+1)(k-i))_+}.$$

Putting everything together, we conclude that $\Psi(X_i) = X_{f(i)}$ in this case too. □

Proof of Lemma 3. Firstly, $f(0) = (k - l)_+ = 0$ because $k < l$. Now let us show by descending induction on i that for $1 \leq i \leq k - 1$, we have $f(i) < i$. Using the fact that $d \geq 1$ and $k \geq 2$, we have $f(k - 1) = (k - d - 1)_+ < k - 1$. Fix $1 \leq i \leq k - 2$. If $f(i) = 0$, we are done, so let us assume that $f(i) > 0$. It is easy to check that f is nondecreasing, thus, $f(i + 1) > 0$. Considering the three cases where i is less than, equal to, or greater than $k - r - 1$, one can prove that $f(i + 1) - f(i) \geq d \geq 1$. Using the induction hypothesis, we obtain $f(i) \leq f(i + 1) - 1 < i + 1 - 1$, which concludes the proof of the second part of the lemma. □

Acknowledgements

We thank Sergey Foss for fruitful discussions. S. R. acknowledges the hospitality of the Mathematical Sciences Research Institute in Berkeley during the ‘Random spatial processes’ program, where this work was initiated.

References

- [1] ALDOUS, D. AND PITMAN, J. (1983). The asymptotic speed and shape of a particle system. In *Probability, Statistics and Analysis*. (London Math. Soc. Lecture Notes Ser. **79**), Cambridge University Press, pp. 1–23.
- [2] FOSS, S. AND KONSTANTOPOULOS, T. (2003). Extended renovation theory and limit theorems for stochastic ordered graphs. *Markov Process. Relat. Fields* **9**, 413–468.
- [3] FOSS, S. AND ZACHARY, S. (2013). Stochastic sequences with a regenerative structure that may depend both on the future and on the past. *Adv. Appl. Prob.* **45**, 1083–1110.
- [4] MALLEIN, B. AND RAMASSAMY, S. (2016). Barak–Erdős graphs and the infinite-bin model. Preprint. Available at <https://arxiv.org/abs/1610.04043>.
- [5] TRAHMAN, A. N. (2011). Modifying the upper bound on the length of minimal synchronizing word. In *Fundamentals of Computation Theory* (Lecture Notes Comput. Sci. **6914**), Springer, Heidelberg, pp. 173–180.