



RESEARCH ARTICLE

Soliton Decomposition of the Box-Ball System

Pablo A. Ferrari¹, Chi Nguyen^{1,2}, Leonardo T. Rolla^{1,3} and Minmin Wang^{1,4}

¹Instituto de Investigaciones Matemáticas Luis A. Santaló, Argentinian National Research Council at the University of Buenos Aires, 1428 Buenos Aires, Argentina.

²Department of Information Technology, Ho Chi Minh City University of Transport, 2 Vo Oanh St., Binh Thanh Dist., Ho Chi Minh City, Vietnam.

³Department of Statistics, University of Warwick, CV4 7AL, United Kingdom.

⁴Department of Mathematics, University of Sussex, BN1 9QH, United Kingdom.


Received: 28 October 2020; **Revised:** 23 June 2021; **Accepted:** 3 July 2021

2020 Mathematics Subject Classification: *Primary* – 37B15; *Secondary* – 60K35, 82C20, 82C70

Keywords: box-ball systems; solitons; invariant measures

Abstract

The box-ball system (BBS) was introduced by Takahashi and Satsuma as a discrete counterpart of the Korteweg-de Vries equation. Both systems exhibit solitons whose shape and speed are conserved after collision with other solitons. We introduce a slot decomposition of ball configurations, each component being an infinite vector describing the number of size k solitons in each k -slot. The dynamics of the components is linear: the k th component moves rigidly at speed k . Let ζ be a translation-invariant family of independent random vectors under a summability condition and η be the ball configuration with components ζ . We show that the law of η is translation invariant and invariant for the BBS. This recipe allows us to construct a large family of invariant measures, including product measures and stationary Markov chains with ball density less than $\frac{1}{2}$. We also show that starting BBS with an ergodic measure, the position of a tagged k -soliton at time t , divided by t converges as $t \rightarrow \infty$ to an effective speed v_k . The vector of speeds satisfies a system of linear equations related with the generalised Gibbs ensemble of conservative laws.



BBS dynamics for independent and identically distributed initial configuration with density 0.25. Time is going down. Straight red lines are deterministic and computed using Theorem 1.2. (High resolution, color online.)

1. Overview

Assume that there is a *box* at each integer $x \in \mathbb{Z}$ and that each box may contain a *ball* or be empty. Denote $\eta \in \{0, 1\}^{\mathbb{Z}}$ a ball configuration, with the convention $\eta(x) := 1$ if there is a ball at x , else $\eta(x) := 0$. Imagine a carrier that traverses \mathbb{Z} from left to right as follows. When visiting box x , the carrier picks a ball if there is one and deposits a ball if the box x is empty and the carrier has at least one ball. Let $T\eta$ be the configuration obtained after the carrier visited all boxes. An example of η , carrier load, and $T\eta$ is as follows:

```

. . . 0 0 0 1 0 1 1 0 0 0 0 1 1 0 1 0 0 0 0 0 0  η
. . . 0 0 0 1 0 1 2 1 0 0 0 1 2 1 2 1 0 0 0  carrier load
. . . 0 0 0 0 1 0 0 1 1 0 0 0 0 1 0 1 1 0 0 0  Tη
    
```

The map T for a general $\eta \in \{0, 1\}^{\mathbb{Z}}$ is defined in (1.2). This dynamics is called the *box-ball system* (BBS) and was introduced by Takahashi and Satsuma [TS90], who proposed an algorithm to identify *solitons* in configurations with a finite number of balls and argued that each soliton identified at time 0 can be tracked at successive iterations of T . For example, in the configuration having exactly three balls at boxes 1, 2, 3 the 3-soliton γ consists of these three occupied boxes and the empty boxes 4, 5, 6. Evolving this configuration by t iterations of T , the new configuration will have a 3-soliton γ^t that is a translation of γ by $3t$. In general, [TS90] observed that a k -soliton always consists of k occupied boxes and k empty boxes. The relative positions can change and be more scattered during collisions with other solitons, but the striking property of the BBS is that such collisions neither create nor destroy solitons. The distance between solitons of the same size is also conserved after collisions.

The main goal of [TS90] was to propose an integrable system with the same behavior as the Korteweg-de Vries equation, KdV, whose solutions include solitons of different sizes that keep shape after collision with other solitons. Then [TTMS96, TM97] argued a way to go from KdV to BBS via ultradiscretisation and tropical geometry; see also [KTZ17, Zuk20]. Further use of Bethe ansatz to study asymptotic behavior of BBS can be found in [IKO04, MIT06, IKT12, KLO18]. The model has also attracted attention in the combinatorics and probability communities. Using a map between solitons and Dick paths in [TTS96], the paper [LLP17] relates soliton sizes to longest decreasing subsequence of restricted permutations. [FG20b] found a new soliton identification that maps to a branch decomposition of the Neveu-Aldous trees of random walks. See [LPS21, Sak14a, Sak14b] for some other combinatorial developments. The paper [HMO01] shows that stationary Markov chains are invariant measures for the Pitman transformation [Pit75], a dynamics equivalent to BBS; see [CKST18, CS19, CS20] for extensions.

This article has three main contributions. Firstly, we discover the following fact: any ball configuration η with ball density less than $\frac{1}{2}$ can be mapped to a family of soliton components, $(\zeta_k)_{k \geq 1}$, called the *slot decomposition* of η ; here $\zeta_k(i)$ represents the number of k -solitons at coordinate $i \in \mathbb{Z}$ of the k th component. The components of a ball configuration evolve linearly under T : component k moves rigidly at speed k . The interaction among components reappears when the ball configuration is reconstructed from the components. This is a delicate hierarchical arrangement where the k th component is *appended* to a certain subset of \mathbb{Z} called *k-slots*, determined by the m -components for $m > k$. The above facts are purely deterministic. It would be interesting to understand the relationship between the rigged configurations in [IKO04] with the slot decomposition.

The BBS can be seen as a dynamical system acting on the set of configurations with density less than $\frac{1}{2}$, and a natural question is about the invariant measures μ defined by $\mu T^{-1} = \mu$. Our second main result states that given a translation-invariant random family ζ of independent vectors satisfying a summability condition, the law of the random ball configuration whose slot decomposition is ζ is translation invariant and invariant for the BBS. Product measures and stationary Markov chains with density less than $\frac{1}{2}$ satisfy those properties, as well as a large family of measures based on soliton weights [FG20a]. We conjecture that the slot decomposition characterises T -invariant probability measures in the sense that if μ is shift-mixing and T -invariant, then its components should be independent and shift-mixing; see Remark 4.8.

Our third main contribution is the characterisation of the *effective soliton speeds* for shift-ergodic initial states, illustrated by the red lines in the figure in the abstract. The result is based on the following rough description of the soliton dynamics. A size k soliton travels at speed k in absence of other solitons and when two solitons collide, the smaller soliton gets delayed and the bigger soliton ‘jumps’ over the slower one. Our proof is based on shift-ergodicity and Palm theory and does not use the slot decomposition. For a nonhomogeneous initial condition, the effective speed equations have been derived by [CS21], who also performed a hydrodynamic limit, using our slot decomposition. The results are analogous to those for the hard rod system [BDS83], where disjoint segments of fixed size of the real line called rods travel ballistically at assigned speeds until collision, when the speeds are interchanged. A *pulse* follows the rod that is travelling at one of the given speeds. Hard rod pulses and BBS solitons have a similar dynamics and, consequently, similar hydrodynamic and effective speed equations. These results belong to the very active area of generalised hydrodynamics of the generalised Gibbs ensemble; see [Spo12, DS17, DYC18, CBS19, KMP20] and references therein.

This article is organised as follows. In Section 1 we state the main results of this article, after giving the definitions needed for the statements. In Section 2 we introduce the slot decomposition of ball configurations, show that the slot decomposition is an injective map (Theorem 1.3) and describe how a configuration can be reconstructed from the components. In Section 3 we show that under the BBS dynamics each component shifts rigidly (Theorems 1.4 and 3.1). In Section 4 we give an explicit construction of T -invariant measures that are shift-invariant (Theorem 1.5). In Section 5 we compile fragments of Palm theory from the literature that play an important role in many of our arguments. In Section 6 we study the asymptotic speed of tagged solitons (Theorems 1.1 and 1.2). We also study the soliton speeds in terms of tagged records (Theorem 6.1). In Section 7 we complete some proofs postponed in previous sections.

1. Preliminaries and results

In this section we describe our main results. We will work mostly with configurations with ball density less than $\frac{1}{2}$. More precisely, let

$$\mathcal{X}_\lambda := \left\{ \eta \in \{0, 1\}^{\mathbb{Z}} : \lim_{y \rightarrow \infty} \frac{1}{y} \sum_{x=-y}^0 \eta(x) = \lim_{y \rightarrow \infty} \frac{1}{y} \sum_{x=0}^y \eta(x) = \lambda \right\},$$

$$\mathcal{X} := \bigcup_{0 < \lambda < \frac{1}{2}} \mathcal{X}_\lambda.$$

In the sequel, a site $x \in \mathbb{Z}$ is often referred to as a *box*. For $\eta \in \mathcal{X}$ we define the set of *records* by

$$R\eta := \left\{ x \in \mathbb{Z} : \sum_{y=z}^x \eta(y) < \sum_{y=z}^x [1 - \eta(y)] \text{ for all } z \leq x \right\}. \tag{1.1}$$

Note that $\eta(x) = 0$ for all $x \in R\eta$. The piece of configuration between two consecutive records forms a finite *excursion*. The operator T is defined by

$$T\eta(x) := \begin{cases} 0, & x \in R\eta, \\ 1 - \eta(x), & \text{otherwise.} \end{cases} \tag{1.2}$$

In other words, the value at records remains 0 and the excursions are flipped. When applied to finite ball configurations, this operator coincides with the operator described in the previous section. We show in Section 2 that if $\lambda < \frac{1}{2}$ then \mathcal{X}_λ is invariant under T .

1.1. Identifying and tracking solitons

Define the *runs* of η as maximal blocks of successive sites where η has a constant value, so that they form a partition of \mathbb{Z} . Assume first that η has a finite number of balls, so it has a finite number of finite runs and two semi-infinite runs of zeros, one to the left and one to the right.

A *k-soliton* is a collection of $2k$ boxes identified by the Takahashi–Satsuma algorithm [TS90] running as follows.

```

Start with a doubly infinite word, where each letter in the word is 0 or 1.
while there are still ones in the word do
    Select the leftmost run in the word whose length is at least as long as the length (denote
    it  $k$ ) of the run preceding it
    Identify a soliton of size  $k$ , or simply  $k$ -soliton, consisting of the first  $k$  letters of this run
    and the letters of the run preceding it
    Remove these  $2k$  letters from the word
end
    
```

Notice that a k -soliton consists of k zeros (possibly nonconsecutive) followed by k ones or vice versa, and letters that do not belong to any soliton are all zero and correspond to the records of η ; see Figure 1.1. For a general $\eta \in \mathcal{X}$, all excursions have a finite number of boxes. To identify the solitons in η , we take each excursion of η , append infinitely many zeros to the left and right of the excursion and then apply the above algorithm to it.

We define the *head* and *tail* of a k -soliton γ as follows: the head $\mathcal{H}(\gamma) = \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k\} \subseteq \mathbb{Z}$ is the set of k boxes with ones in γ and the tail $\mathcal{T}(\gamma) = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_k\} \subseteq \mathbb{Z}$ is the set of k boxes with zeroes in γ . Namely, $\mathcal{H}(\gamma) \cup \mathcal{T}(\gamma)$ is the set of boxes that are removed when executing the previous algorithm on η . Let $\Gamma_k \eta$ be the set of k -solitons of a ball configuration $\eta \in \mathcal{X}$. The following is proved in Section 7.

Proposition 1.3. *For any $\eta \in \mathcal{X}$ and $A \subseteq \mathbb{Z}$, there is a k -soliton $\gamma \in \Gamma_k \eta$ with tail $\mathcal{T}(\gamma) = A$ if and only if there is a k -soliton $\gamma^1 \in \Gamma_k(T\eta)$ with head $\mathcal{H}(\gamma^1) = A$.*

By the above proposition, we can track each k -soliton γ in the evolution of η . For each k -soliton $\gamma \in \Gamma_k \eta$, call $(\gamma^t)_{t \geq 0}$ the trajectory satisfying $\gamma^0 = \gamma$, $\gamma^t \in \Gamma_k T^t \eta$ and

$$\mathcal{H}(\gamma^{t+1}) = \mathcal{T}(\gamma^t). \tag{1.4}$$

1.2. Soliton nesting and motion

As shown in Figure 1.1, solitons can be nested inside larger solitons. As it turns out, they are nested in a hierarchical way. Moreover, solitons only move to the right, and they are only free to move when they




Figure 1.1. Applying the Takahashi–Satsuma algorithm to a sample configuration. Dots represent records. On the left we have the resulting word after successive iterations. Identified solitons are shown in bold once and then with a color corresponding to their size. The algorithm is applied to each excursion separately, so the rightmost 1-soliton in the picture is ignored by this instance of the procedure. (color online)

The soliton speeds v_k are finite, positive, increasing in k and satisfy the system

$$v_k = k + \sum_{m < k} 2m\bar{\rho}_m(v_k - v_m) - \sum_{m > k} 2k\bar{\rho}_m(v_m - v_k). \tag{1.10}$$

System (1.10) comes from the following. When a k -soliton is isolated, it advances by k units, and when it encounters an m soliton, the encounter causes it to advance $2m$ extra units if $m < k$ or stay put for two units of time if $m > k$. The term $\bar{\rho}_m|v_k - v_m|$ gives the frequency of such encounters as seen from a k -soliton.

When $\rho_k = 0$, the soliton speed v_k does not come from (1.9) but in principle formally from (1.10). There is still an interpretation for v_k in terms of the dynamics, as discussed in Subsection 6.5.

When $\rho_k > 0$ for finitely many k , the system has a unique solution [CS21, Lemma 5.1]. We believe it is unique in general, but we have been unable to prove it. The following partial result is proved in Section 7.

Proposition 1.11. *If $\sum_m m\bar{\rho}_m < \frac{1}{4}$, then the nonnegative solution to (1.10) is unique.*

As a side remark, $\sum_m m\bar{\rho}_m = \lambda$; that is, the mean number of occupied boxes per unit space. Uniqueness has not been proved to hold in general, and we show that $(v_k)_k$ is determined by the vector $(\bar{\rho}_k)_k$ under a stronger assumption in terms of soliton components (described in Subsection 1.5).

Theorem 1.2. *If, when conditioned on having a record at $x = 0$, μ has independent soliton components and each component is independent and identically distributed (i.i.d.), then the soliton speeds $(v_k)_k$ in (1.9) are also given by the unique solution to*

$$\begin{aligned} w_k &= 1 + \sum_{m > k} 2(m - k)\rho_m, & \rho_k &= \alpha_k w_k, \\ s_k &= k + \sum_{m < k} 2(k - m)s_m \alpha_m, & v_k &= \frac{s_k}{w_k}, \end{aligned} \tag{1.12}$$

and, in particular, they are determined by $(\rho_k)_k$.

In system (1.12), w_k is the density of k -slots per excursion (see Subsection 1.4 for the definition of k -slot), α_k is the density of k -solitons per k -slot, s_k is the average size of the head of a k -soliton, $k - m$ is the number of m -slots in the head of a k -soliton and the factor $\frac{1}{w_k}$ is the probability that a typical k -soliton is free to move (see Section 6 for details).

The proof of (1.12) uses independence properties of the components for an explicit computation of the mean jump size of a typical k -soliton γ in one iteration. By ergodicity, the mean jump size is v_k , the limit of $x(\gamma^t)/t$, as shown in Subsection 6.1.

In the setup of Theorem 1.1, we cannot compute the mean jump size explicitly. However, if we assume that the solution to (1.10) is indeed unique, then by taking an initial measure with independent components and the same vector $(\rho_k)_k$, we see that the vector $(v_k)_k$ must be given by (1.12). So if the solution to (1.10) is indeed unique (as conjectured), the independence assumption in Theorem 1.2 can be waived.

In practice, the soliton speeds $(v_k)_k$ can be computed by truncating ρ (replace ρ_k by 0 for large k) and solving these finite recursions for w , α , s and finally v . When the initial ball configuration consists of i.i.d. Bernoulli random variables, one can find α_k explicitly in terms of the density λ by computing partition functions [FG20a], substitute the equation for ρ into that for w and then compute s and v . Using this and the above theorem, we have found the asymptotic speeds of the solitons for the simulations shown in Figures 1.3 and 1.4 as well as that on the first page.




Figure 1.3. Simulation for an i.i.d. configuration with density 0.15. The transparent red lines have deterministic slopes computed by Theorem 1.2, which have been manually shifted so that they would overlay a soliton. This window covers 2000 sites and 140 time steps going downwards and has been stretched vertically by a factor of 5. The figure on the first page is the same except for the density. (high resolution, color online)




Figure 1.4. Simulation for $(\rho_k)_k = (.006, .005, .1, .003, 0, 0, 0, \dots)$. The initial configuration was obtained by first appending one k -soliton with probability ρ_k after each record and then applying T a number of times in order to mix. As in Figure 1.3, it is a 2000 x 140 window stretched by 5, and red lines are deterministic. (high resolution, color online)

1.4. Slot decomposition

Recall that a k -soliton has a head and a tail, each one consisting of k (possibly nonconsecutive) sites. We say that the j th box of the head or tail of an m -soliton is a k -slot for all $k < j$ and that a record is a k -slot for every k . Roughly speaking, the k -slots are the places where k -solitons can be appended; see Subsection 2.1 for precise definitions and examples.

The set of configurations with a record at the origin is defined by

$$\widehat{\mathcal{X}} := \{\eta \in \mathcal{X} : 0 \in R\eta\}.$$

Assume that $\eta \in \widehat{\mathcal{X}}$. Enumerate the k -slots from left to right such that the 0th k -slot is at position $s_k(\eta, 0) = 0$, and let $s_k(\eta, i)$ denote the position of the i th k -slot for $i \in \mathbb{Z}$. We say that a k -soliton γ is appended to the i th k -slot if its head and tail are contained between $s_k(\eta, i)$ and $s_k(\eta, i + 1)$. Define the

k th component of η as the configuration $M_k\eta$ given by $M_k\eta(i) :=$ number of k -solitons appended to the i th k -slot. Define

$$\mathcal{M} := \left\{ \zeta = (\zeta_k)_{k \geq 1}, \zeta_k \in (\mathbb{Z}_+)^{\mathbb{Z}} : \sum_k \zeta_k(i) < \infty, \text{ for all } i \in \mathbb{Z} \right\}.$$

The next result, proved in Subsection 2.2, says that we can recover η from its components $(M_k\eta)_k$.

Theorem 1.3. *For $\eta \in \widehat{\mathcal{X}}$, we have $(M_k\eta)_k \in \mathcal{M}$. Moreover, the map $M : \eta \in \widehat{\mathcal{X}} \mapsto (M_k\eta)_k$ is invertible.*

The fundamental property of this decomposition is that it makes the BBS dynamics linear. We show a simple case, deferring the full statement until Theorem 3.1 in Section 3.

Theorem 1.4. *Suppose $\eta \equiv 0$ on \mathbb{Z}_- and η has infinitely many records on \mathbb{Z}_+ . Then for all $k \geq 1$,*

$$M_k T \eta = \theta^{-k} M_k \eta.$$

So, when we see a configuration η through its components, the k th component is displaced by k units, without interacting with the other components. In the general case, there will be larger solitons arriving from the left, which will disrupt the enumeration of k -slots, making the statement substantially more involved. See the details in Section 3.

1.5. Invariant measures

Using the slot decomposition and the reconstruction map $(M_k\eta)_k \mapsto \eta$ that will be described in detail in Subsection 2.2, we get the following.

Theorem 1.5. *Let $\zeta = (\zeta_k)_{k \geq 1}$ be independent random elements of $(\mathbb{Z}_+)^{\mathbb{Z}}$ with shift-invariant distributions satisfying $\sum_k k E[\zeta_k(0)] < \infty$ and $P(\sum_{k,i} \zeta_k(i) > 0) = 1$. Then there exists a unique shift-invariant probability μ on \mathcal{X} such that $M_k\eta \stackrel{d}{=} \zeta_k$ when η is distributed with μ conditioned on $\widehat{\mathcal{X}}$. This measure μ is T -invariant. If, moreover, $(\zeta_k(i))_{i \in \mathbb{Z}}$ is i.i.d. for each k , then μ is also shift-ergodic.*

The above theorem says that the family of invariant measures for this dynamics is at least as large as the family of sequences of states of k -soliton configurations. In particular, given a sequence $(\alpha_k)_k$ specifying the density of k -solitons in the k th component, we can construct an infinite number of mutually singular shift-invariant and T -invariant laws μ on \mathcal{X} , all having the same specified component densities.

The extra assumption needed in Theorem 1.2 is that μ be of the above form; that is, conditioning on $\widehat{\mathcal{X}}$, each k th component is i.i.d. and they are independent over k . In this case, we can also study the speed of tagged records and the speed of solitons measured in terms of tagged records; see Subsection 6.5. As pointed out above, this condition should not be necessary for Theorem 1.2 to hold.

We should also note that the converse of Theorem 1.5 is false. In particular, there exist invariant measures that are not constructed from independent components; see Examples 4.6 and 4.7 as well as Remark 4.8 in Subsection 4.1.

2. Slot decomposition

Let $\xi = \xi[\eta]$ be a walk on \mathbb{Z} that jumps one unit up at x when there is a ball at x and jumps one unit down when box x is empty. That is,

$$\xi(x) - \xi(x - 1) = 2\eta(x) - 1.$$

Note that for each η , such a walk $\xi[\eta]$ is not unique and only defined up to a vertical shift. We define records for a walk ξ in the usual sense; that is, we say that x is a record for ξ if $\xi(z) > \xi(x)$ for all $z < x$.




Figure 2.1. Time evolution of a walk under seven iterations of T . This example has four solitons, of size 7, 5, 3 and 1. Different colors are used to highlight their conservation. To facilitate view, we have shifted the walk at time t by t units down. (color online)

Let \mathcal{W}_λ be the set of all simple walks ξ such that $\eta[\xi] \in \mathcal{X}_\lambda$, which is the set of walks ξ such that


$$\lim_{x \rightarrow \pm\infty} \frac{\xi(x)}{x} = 2\lambda - 1, \tag{2.1}$$

and $\mathcal{W} = \cup_{0 < \lambda < 1/2} \mathcal{W}_\lambda$. Then every $\xi \in \mathcal{W}$ satisfies $\min_{y \leq x} \xi(y) \in \mathbb{Z}$ for all $x \in \mathbb{Z}$, and we define

$$T\xi(x) := 2 \min_{y \leq x} \xi(y) - \xi(x) = \left[\min_{y \leq x} \xi(y) \right] - \left[\xi(x) - \min_{y \leq x} \xi(y) \right]. \tag{2.2}$$

This amounts to reflecting the walk ξ with respect to its running minimum. It is worth mentioning that the above transformation for Brownian motion was studied by Pitman [Pit75]. The way that T acts on \mathcal{W} is illustrated with an example in Figure 2.1.

One can see ξ as a lift of η that includes an arbitrary choice of vertical shift or, equivalently, an arbitrary labelling of records in increasing order. Conversely, $\eta[\xi]$ is unambiguously defined by $\eta(x) = \frac{1 + \xi(x) - \xi(x-1)}{2}$. Consider the following diagram:



Notice that the above definition of record coincides with the one given in (1.1) and (1.2) is equivalent to (2.2). Therefore, this diagram commutes except that the lifting \mathcal{L} misses uniqueness and the projection \mathcal{P} cancels such nonuniqueness. They are analogous to the derivative and indefinite integral where the




Figure 2.2. A red 3-soliton can be appended to a blue 5-soliton in $2 \times (5 - 3) = 4$ different places, represented by blue dots. It is also possible to append it to records, represented by black dots. Attempting to insert a 3-soliton at a site not marked by a dot will result in erroneous soliton identification. For instance, the 3-soliton in the middle bottom plot should actually start three sites earlier, and in the right bottom plot we should have a 2- and 6-soliton instead of 3 and 5. The green crosses indicate that the coloring is inconsistent with the procedure shown in Figure 1.1. (color online)

latter comes with an indeterminate additive constant. If a property is insensitive to the choice of the lift $\xi[\eta]$, then it is in fact a property of η , even if is described in terms of ξ . For instance, for $\lambda < \frac{1}{2}$, T -invariance of \mathcal{X}_λ is equivalent to the T -invariance of \mathcal{W}_λ , which follows immediately from (2.1) and (2.2). Note that properties of η always translate to ξ ; for instance, $\Gamma_m \xi$ means simply $\Gamma_m \eta[\xi]$, etc. However, some of the objects considered in this section do depend on the lift ξ .

2.1. Slots and components

We now describe how solitons can be nested inside each other via what we call *slots*. Intuitively, the idea of k -slot is that it marks a place where a k -soliton can be inserted without interfering with the rest of the configuration in terms of the Takahashi–Satsuma algorithm; see Figure 2.2.

Let $\gamma \in \Gamma_k(\xi)$ be a k -soliton for the walk representation ξ . We label the sites in the head (respectively in the tail) of γ from left to right: $\mathcal{H}(\gamma) = \{\mathcal{H}_1(\gamma), \dots, \mathcal{H}_k(\gamma)\}$ (respectively $\mathcal{T}(\gamma) = \{\mathcal{T}_1(\gamma), \dots, \mathcal{T}_k(\gamma)\}$). The *slot configuration* $S\xi : \mathbb{Z} \rightarrow \{0, 1, 2, \dots\} \cup \{\infty\}$ is defined by

$$S\xi(x) := \begin{cases} m - 1, & \text{if } x = \mathcal{T}_m(\gamma) \text{ or } \mathcal{H}_m(\gamma), \text{ for } \gamma \in \Gamma_k \xi \text{ and } k \geq m, \\ \infty, & \text{if } x \text{ is a record for } \xi. \end{cases}$$

For each $k \geq 1$ we say that x is a k -slot for ξ if $S\xi(x) \geq k$.

Note that a record is a k -slot for all k , and an m -soliton contains a number $2m - 2k$ of k -slots; see Figure 2.3. Because every $\xi \in \mathcal{W}$ has infinitely many records, it also has infinitely many k -slots.

For $j \in \mathbb{Z}$, the position of the record at level $-j$ will be called *record j* and denoted as

$$r(\xi, j) := \min\{x \in \mathbb{Z} : \xi(x) = -j\}.$$

This is the leftmost site where the walk ξ takes the value $-j$. If $\xi \in \mathcal{W}$, we have $r(\xi, j) \in \mathbb{Z}$ well defined for all $j \in \mathbb{Z}$. The site $r(\xi, 0)$ will play a central role in the sequel. Note that $T\xi(x) \leq \xi(x)$ by (2.2), so $r(T\xi, j) \leq r(\xi, j)$ for all j .

We label all k -slots in increasing order: $\dots < s_k(\xi, -1) < s_k(\xi, 0) < s_k(\xi, 1) < \dots$, where $s_k(\xi, 0) := r(\xi, 0)$. The set of k -slots of ξ is denoted $S_k \xi := \{s_k(\xi, i) : i \in \mathbb{Z}\}$. We then say that a k -soliton γ is *appended* to the i th k -slot $s_k(\xi, i)$ if $\gamma \cap [s_k(\xi, i), s_k(\xi, i + 1) - 1] \neq \emptyset$. Observe that if that is the case, we necessarily have $\gamma \subseteq [s_k(\xi, i) + 1, s_k(\xi, i + 1) - 1]$, as a consequence of Lemma 1.5 and the fact that k -slots can only be records or sites of solitons with larger size. It follows that each k -soliton in ξ is appended to a unique k -slot. On the other hand, it is possible to have multiple k -solitons appended to a single k -slot. Finally, we let $M_k \xi(i)$ be the number of k -solitons appended to the i th k -slot and call

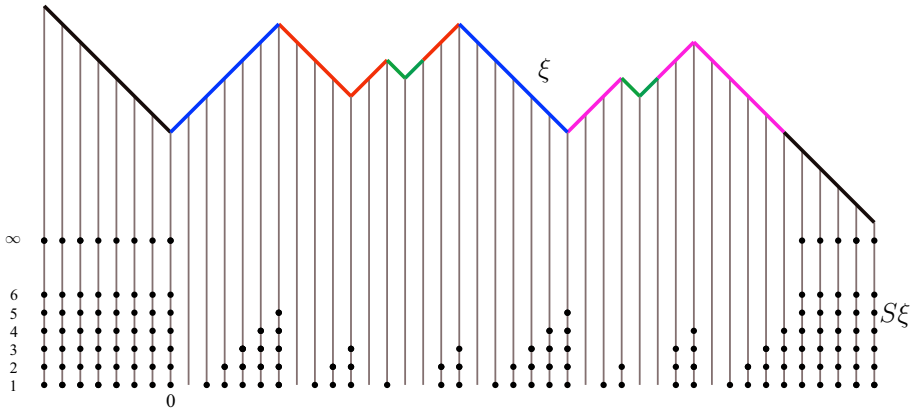


Figure 2.3. Slot configuration of a walk ξ . Different colors correspond to different solitons; records are painted in black. For each site, the number of dots below it indicates its level in slots: k dots indicate a k -slot. (color online)




Figure 2.4. An illustration of how the solitons are nested inside bigger solitons via slots, in the same sample configuration as in Figure 2.3. Solitons are represented by squares and slots by circles. For each $k \geq 1$, each slot with index $m \geq k$ is a k -slot. We say it is the n th k -slot, where n is determined by counting how many k -slots appear before it in the depth-first order, and the counting starts from the 0th k -slot present at record 0.

$M_k \xi = (M_k \xi(i))_{i \in \mathbb{Z}}$ the k th component of ξ . For instance, in the example of Figure 2.3, we have $M_6 \xi(0) = 1$, $M_5 \xi(2) = 1$, $M_4 \xi(2) = 1$, $M_1 \xi(9) = 1$, $M_1 \xi(18) = 1$ and $M_k \xi(i) = 0$ otherwise. See also Figure 2.4 for an illustration on how the solitons are nested inside each other via slots.

2.2. Reconstructing the configuration from the components

Here we prove Theorem 1.3. In this subsection only, we will work with a larger set of configurations than \mathcal{W} . Let \mathcal{W}_* be the set of walks ξ such that $r(\xi, j)$ is well defined for all $j \in \mathbb{Z}$. Even though \mathcal{W}_* is not preserved by T , the Takahashi–Satsuma algorithm still applies, because all of the excursions of $\xi \in \mathcal{W}_*$ are finite. So the k th component M_k described in Subsection 2.1 is also defined for all $\xi \in \mathcal{W}_*$.

On the other hand, because the slot decomposition $\xi \mapsto M\xi$ is insensitive to horizontal shifts, it is not possible to determine ξ knowing $(M_k\xi)_{k \geq 1}$. So we introduce

$$\widehat{\mathcal{W}}_* := \{\xi \in \mathcal{W}_* : r(\xi, 0) = 0\}.$$

(As a side remark, \mathcal{W}_* is the set of simple walks ξ such that $\lim_{x \rightarrow -\infty} \xi(x) = +\infty$ and $\liminf_{x \rightarrow +\infty} \xi(x) = -\infty$. We can build a configuration $\xi \in \mathcal{W}_*$ with $T\xi \notin \mathcal{W}_*$ by appending a soliton of height 2^n between records $-n$ and $-n + 1$, so that $\liminf_{x \rightarrow -\infty} T\xi(x) = -\infty$.)

Note that, unlike the lift $\xi[\eta]$ from \mathcal{X} to \mathcal{W} , which was not unique, for η in $\widehat{\mathcal{X}}$ there is a unique lift $\xi[\eta]$ that is in $\widehat{\mathcal{W}}_*$. We denote this unique lift by $\xi^\circ[\eta]$. So the maps $\eta \mapsto M_k\eta$ as seen in Theorem 1.3 are given by $M_k\eta = M_k\xi^\circ[\eta]$. But to remain consistent with the previous subsections, we continue to work with ξ instead of η .

Denote by M the map $\xi \mapsto M\xi := (M_k\xi)_{k \geq 1}$. We now show that $M : \widehat{\mathcal{W}}_* \rightarrow \mathcal{M}$ is invertible.

The *height* of the excursion between record j and record $j + 1$, denoted as $m(j)$, is defined as 0 if the excursion is empty or as the largest k such that a k -soliton is contained in the excursion. Denote also $i_k(j)$ the label of the k -slot located at record j . Because k -slots can only be created by solitons of larger sizes, we note that

$$m(j) = \min\{k \geq 0 : M_{k'}\xi(i_{k'}(j)) = 0 \text{ for all } k' > k\} \in \{0, 1, 2, \dots\}.$$

Both $m(j)$ and $i_k(j)$ depend on ξ but we omit it in the notation. Denote the slot decomposition $M\xi$ of ξ by ζ ; that is,

$$\zeta = (\zeta_k)_k = M\xi.$$

Note that each $\xi \in \mathcal{W}_*$ and $j \in \mathbb{Z}$, $i_k(j) = r(\xi, j)$ for all $k \geq \max\{m(0), \dots, m(j)\}$. Because $m(j)$ is finite, we have that

$$\text{for each } i \in \mathbb{Z}, \zeta_k(i) = 0 \text{ for all large } k. \tag{2.3}$$

Namely, $\zeta \in \mathcal{M}$. Conversely, suppose that $\zeta = (\zeta_k)_{k \geq 1} \in \mathcal{M}$, so, in particular ζ satisfies (2.3). We first give an algorithm to reconstruct the excursion of ξ between records 0 and 1. To that end, we introduce the following notation: denote the number of k -slots in the excursion ε between successive records $y_1 < y_2$ by

$$n_k(\varepsilon) := 1 + |S_k\xi \cap \{y_1 + 1, \dots, y_2 - 1\}|, \tag{2.4}$$

where the term 1 refers to the record y_1 preceding ε and the second term counts the number of k -slots belonging to m -solitons of ε with $m > k$. Here is the algorithm:

Let ε be the empty configuration.

Let $m := \min\{k \geq 0 : \zeta_{k'}(0) = 0 \text{ for all } k' > k\}$

for $k = m, m - 1, \dots, 2, 1$ **do**

Let $n_k := n_k(\varepsilon) = \#\{x \in S_k\varepsilon : r(\varepsilon, 0) \leq x < r(\varepsilon, 1)\}$, as in (2.4)

for $i = 0, 1, \dots, n_k - 1$ **do**

Insert a number $\zeta_k(i)$ of k -solitons in the i th k -slot of ε ; that is, to the right of site $x = s_k(\varepsilon, i)$; boxes to the right of x are shifted further right in order to accommodate the insertion of these k -solitons

This produces an updated configuration ε

end

end

Reconstruction of ε^0




Figure 2.5. Reconstruction algorithm for a single excursion. This example is obtained using the field ζ shown in Figure 2.6.

Note that m is well defined by (2.3). In case $m = 0$, the algorithm produces an empty configuration. Let us also note that when inserting solitons in the above algorithm, there is only one way to do it consistently with the soliton decomposition: if a soliton is inserted to the right of the site x with $\eta(x) = 0$, then it has its head on the left and tail on the right (i.e., $11 \cdots 100 \cdots 0$); otherwise, it is inserted with its tail on the left and head on the right (i.e., $00 \cdots 011 \cdots 1$). The procedure is illustrated in Figure 2.5. Call ε^0 the excursion between record 0 and record 1 obtained from ζ as just described. Construct ε^1 , the excursion between records 1 and 2, using the same algorithm but with the data $\zeta^1 = (\zeta_k^1)_{k \geq 1}$, where each component is given by

$$\zeta_k^1 = (\zeta_k(n_k + i))_{i \geq 0},$$

which consists of the entries of ζ with nonnegative indices i not used in the reconstruction of ε^0 . Note that ζ^1 also satisfies (2.3). Iterate this procedure to construct an infinite sequence of excursions $(\varepsilon^j)_{j=0,1,2,\dots}$. See Figure 2.7.

To reconstruct the configuration to the left of record 0 – that is, to obtain the excursions ε^j with negative j – we use an analogous algorithm that uses the entries of ζ with i -indices starting at -1 and moving left instead of starting at 0 and moving right. First take $\zeta^{-1} = (\zeta_k^{-1})_{k \geq 1}$ where each component is given by $\zeta_k^{-1} = (\zeta_k(i))_{i < 0}$ and use ζ^{-1} to construct ε^{-1} . Then define $\zeta^{-2} = (\zeta_k^{-2})_{k \geq 1}$ where each component is given by $\zeta_k^{-2} = (\zeta_k(i - n_k))_{i < 0}$ and use it to construct ε^{-2} . Iterate this procedure to construct an infinite sequence of excursions $(\varepsilon^j)_{j=-1,-2,\dots}$.

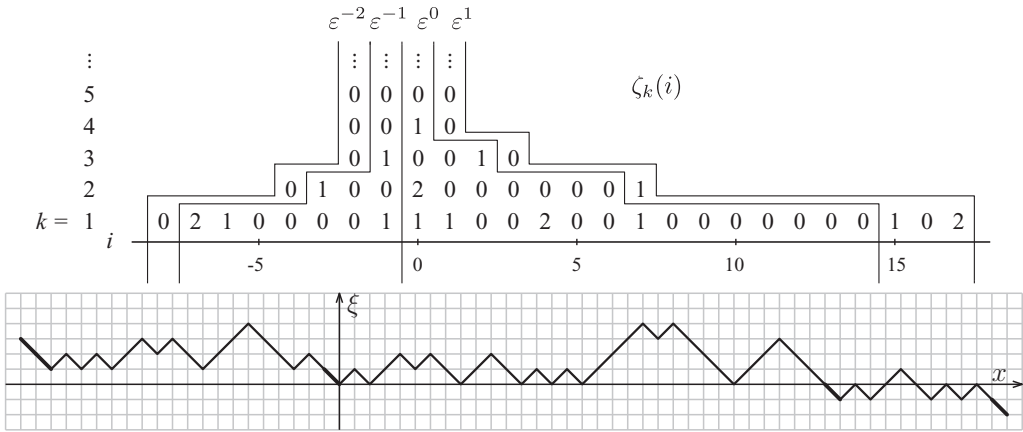


Figure 2.6. Reconstruction of ξ from ζ . In the lower part we show records -2 to 2 in boldface and the excursions between them. Above we show the parts of the field ζ used in the reconstruction of $\varepsilon^{-2}, \varepsilon^{-1}, \varepsilon^0, \varepsilon^1$. Reconstruction of ε^0 was shown in Figure 2.5 and $\varepsilon^1, \varepsilon^{-1}, \varepsilon^{-2}$ is shown in Figure 2.7.

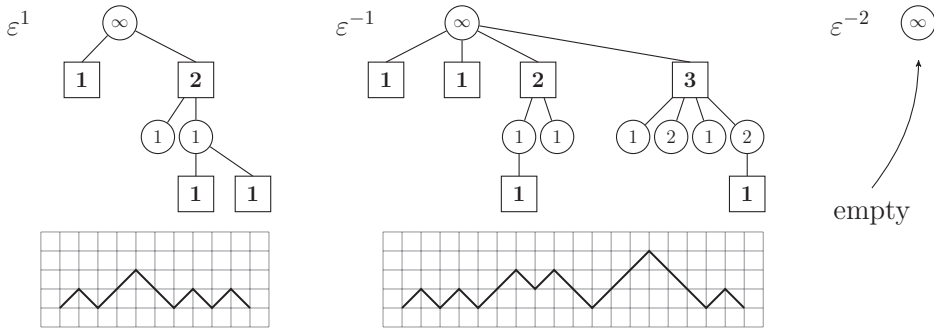


Figure 2.7. Reconstruction algorithm for other excursions. The procedure is the same as in Figure 2.5 but all of the intermediate steps are omitted.

Put record 0 at the origin and concatenate the excursions with one record between each pair of consecutive excursions. This yields a walk denoted as ξ^* , shown in Figure 2.6. Note that in ξ^* , all excursions are finite and therefore $r(\xi, j)$ is finite for all $j \in \mathbb{Z}$. Namely, $\xi^* \in \widehat{\mathcal{W}}_*$.

Call $M^{-1} : \zeta \mapsto \xi^*$ the resulting transformation. We claim that M^{-1} is the inverse map of M ; that is, $M^{-1}M\xi = \xi$ for $\xi \in \widehat{\mathcal{W}}_*$ and $MM^{-1}\zeta = \zeta$ for $\zeta \in \mathcal{M}$. The second identity follows from the fact that $M_k\xi^* = \zeta_k$ for each k . Now let ε be an excursion of ξ . For $k \geq 0$, denote by $\varepsilon_{[k]}$ the ball configuration obtained by removing the boxes belonging to an ℓ -soliton with $\ell \leq k$. Then $\varepsilon_{[0]} = \varepsilon$ and $\varepsilon_{[k]}$ is the empty excursion for k sufficiently large. Now observe from the previous definitions that $M_m\varepsilon_{[k]} = M_m\varepsilon$ for all $m > k$, because m -slots are only created by solitons of sizes larger than m . So the reconstruction algorithm correctly finds $\varepsilon_{[k]}$ from $M_k\varepsilon$ and $\varepsilon_{[k+1]}$ and hence it correctly finds ε .

This shows that $M : \widehat{\mathcal{W}}_* \rightarrow \mathcal{M}$ is invertible, and so is its restriction to a subset $M : \widehat{\mathcal{W}} \rightarrow M(\widehat{\mathcal{W}}) \subseteq \mathcal{M}$, where $\widehat{\mathcal{W}} := \{\xi[\eta] \in \widehat{\mathcal{W}}_* : \eta \in \widehat{\mathcal{X}}\}$. This proves Theorem 1.3.

Let us also point out that the image set $M(\widehat{\mathcal{W}})$ is not simple to characterise. However, as we will see below, if we sample ζ with an appropriate measure, the resulting (random) element $M^{-1}\zeta$ does belong a.s. to $\widehat{\mathcal{W}}$.

3. Evolution of components

Here we prove a stronger version of Theorem 1.4. Recall that we can track a tagged soliton γ after t iterations of T by (1.4). In order to describe how the BBS dynamics act on the components, we introduce the flows of solitons as follows. First, let

$$J_m^t \xi := \#\{\gamma \in \Gamma_m \xi : \gamma \subseteq (-\infty, r(\xi, 0)) \text{ and } \gamma^t \subseteq [r(T^t \xi, 0), \infty)\}.$$

In words, $J_m^t \xi$ counts the number of m -solitons that were to the left of record 0 initially and are found to be on the right of record 0 at time t . We say that such solitons *cross record 0*. Note that this is not the same as ‘crossing the origin’ because record 0 is itself moving left.

Note that $J_m^1 \xi = 0$ for all large m . Indeed, using Proposition 1.3 one can see that no more than $r(\xi, 0) - r(T\xi, 0)$ solitons can cross record 0 at time 1, so $\sum_m J_m^1 \xi < \infty$. By the same argument, $\sum_m J_m^2 \xi < \infty$, and so on.

We now define an observable $o_k^t(\xi)$ that counts the flow of k -slots through record 0 after t iterations of T . Because each m -soliton crossing record 0 from left to right carries $2(m - k)$ k -slots, we define

$$o_k^t(\xi) := \sum_{m>k} 2(m - k) J_m^t \xi < \infty. \tag{3.1}$$

Using the observable $o_k^t(\xi)$, we define the *tagged 0th k -slot at time t* by

$$s_k^t(\xi, 0) := s_k(T^t \xi, o_k^t(\xi)),$$

which is the position of the o_k^t th k -slot counting from record 0 of $T^t \xi$. More generally, the *tagged i th k -slot at time t* is defined as

$$s_k^t(\xi, i) := s_k(T^t \xi, o_k^t(\xi) + i).$$

We now state one of the central results of this article. In some sense, it shows that the action of T on the component configuration $(M_k \xi)_{k \in \mathbb{N}}$ has a simple form, analogous to a Jordan form or upper triangular form for a linear map. The k th component is shifted by k units, plus possibly a number resulting from the re-indexing of k -slots caused by larger solitons crossing record 0.

Theorem 3.1. *Let $\xi \in \mathcal{W}$ and $t \in \{1, 2, 3, \dots\}$. Then, for any k -soliton γ of ξ ,*

$$\#\{i \in \mathbb{Z} : \gamma \subseteq (-\infty, s_k(\xi, i)) \text{ and } \gamma^t \subseteq [s_k^t(\xi, i), \infty)\} = kt;$$

that is, between times 0 and t a tagged k -soliton moves kt units to the right in terms of tagged k -slots or, equivalently, the right-to-left flow of tagged k -slots through a tagged k -soliton is exactly kt . For each $k \in \mathbb{N}$, the k -soliton component of $T^t \xi$ is a shift of the k -soliton component of ξ :

$$M_k T^t \xi(i) = M_k \xi(i - o_k^t(\xi) - kt),$$

where $o_k^t(\xi)$ is defined in (3.1). Moreover, the offset $o_k^t(\xi)$ is determined by $(M_m \xi)_{m>k}$, so that $M_k T^t \xi$ is a function of $(M_m \xi)_{m \geq k}$.

Another way to interpret the theorem is based on the notion of k -bearer, which we introduce now. Let $\pi \in S_k \xi \subseteq \mathbb{Z}$. Then $\pi = s_k(\xi, j)$ for some j , and we define the corresponding k -bearer at time t to be

$$\pi^{k,t} := s_k(T^t \xi, o_k^t(\xi) + kt + j). \tag{3.2}$$

An immediate consequence of Theorem 3.1 is as follows: a k -soliton γ is appended to the k -slot located at $\pi = s_k(\xi, j)$ if and only if γ^t is appended to the k -slot located at $\pi^{k,t}$. In particular, because the definition of a tagged soliton depends only on $\eta[\xi]$ according to Proposition 1.3, this implies that $\pi^{k,t}$

also only depends on $\eta[\xi]$. We note that the difference between the tagged slots introduced before and bearers introduced just now is the factor of kt related to the motion of k -solitons. So a tagged k -soliton crosses k tagged k -slots per unit time, whereas it does not cross k -bearers (in fact, it just follows one of them). As a side remark, the notion of k -bearer allows us to define the soliton speed v_k even if $\rho_k = 0$ by replacing $x(\gamma^t)$ with $\pi^{k,t}$ in (1.9).

Proof of Theorem 3.1. We start by showing how the second statement follows from the first one. Because every k -soliton crosses exactly k tagged k -slots at each step, the number of k -slots between any pair of tagged k -solitons is conserved by T . Hence, the k th component as seen from the tagged 0th k -slot just shifts k k -slots per unit time, and the term $o_k^t(\xi)$ accounts for the relabeling of k -slots caused by bigger solitons crossing record 0.

For the proof of the first statement, it suffices to consider the case $t = 1$. Let us first assume that there is only a finite number of balls in $\xi[\eta]$. Without loss of generality, we can also assume $r(\xi, 0) = 0$. For $x \in \mathbb{Z}_+$, let us define $D_k(\xi, x) = \#S_k T\xi \cap [0, x] - \#S_k \xi \cap [0, x]$, namely, the change in the number of k -slots found in $[0, x]$ after one iteration of T . Let us recall the notation $\mathcal{H}_i(\gamma)$, which stands for the i th site of the head of a soliton γ . We introduce the set $\mathcal{D}_k(\xi, x) = \{z \in [0, x] : \exists \gamma \text{ s.t. } z = \mathcal{H}_i(\gamma), 1 \leq i \leq k \text{ and } \mathcal{T}_i(\gamma^1) > x\}$. Let us observe that if such a soliton γ exists, then necessarily $x(\gamma^1) > x(\gamma)$; otherwise, γ merely swaps its head and tail, and we will have $\mathcal{T}_i(\gamma^1) = \mathcal{H}_i(\gamma)$ in that case. Therefore, $\#\mathcal{D}_k(\xi, x)$ counts the number of slots in $[0, x]$ that have order $< k$ and belong to the heads of solitons that move to the right of x after one iteration. See Figure 3.1 for an example. We claim that for all $x \in \mathbb{Z}_+$ and all $k \geq 1$,

$$D_k(\xi, x) = \#\mathcal{D}_k(\xi, x). \tag{3.3}$$

To prove this, let us first suppose that η consists of one soliton, say, at the sites $[0, 2m]$. Then the slot configuration $ST\xi$ is obtained from $S\xi$ by swapping the labels of the sites in $[0, m - 1]$ with those in $[2m, 3m - 1]$. One readily checks that in this case $D_k(\xi, x)$ is tent shaped: linear with slope 1 on $[0, k \wedge m - 1]$, constant on $[k \wedge m - 1, 2m]$ and null on $[2m + m \wedge k, \infty)$. This can be similarly checked for $\mathcal{D}_k(\xi, x)$ and the claimed identity holds in this case.

For a general configuration ξ , the sites in $\mathcal{D}_k(\xi, x)$ do not necessarily become records in $T\xi$, because they can be claimed by other solitons. However, we will see that even if this happens, it will only affect where the ‘new’ k -slots are located but not the quantity $D_k(\xi, x)$. Proceeding with the proof of (3.3), let us suppose that it holds for configurations with up to $p \geq 1$ solitons and let $\xi[\eta]$ be a configuration with $p + 1$ solitons. Let $m \geq 1$ be the smallest size of the solitons in ξ and let γ be its leftmost m -soliton. We note that $\mathcal{H}(\gamma)$ and $\mathcal{T}(\gamma)$ is back-to-back in ξ (i.e., no other solitons lodged inside γ) as a consequence of Lemma 1.5. Then we can write $I(\gamma) = \mathcal{H}(\gamma) \cup \mathcal{T}(\gamma) = [a, a + 2m]$ with $a = x(\gamma)$. Let us introduce $\tilde{\xi}(x) = \xi(x + 2m)$ for all $x \geq a$ and $\tilde{\xi}(x) = \xi(x)$ otherwise. Namely, $\tilde{\xi}$ is the configuration obtained by removing γ . By noting that the operator T consists in flipping portions of ξ between consecutive records, we deduce that

$$T\xi(x) = T\tilde{\xi}(a - 1) + \sum_{a \leq y \leq x} (1 - 2\eta(y)), \quad \text{if } x \in [a, a + 2m),$$

and $T\xi(x) = T\tilde{\xi}(x - 2m)$ if $x \geq a + 2m$, and $T\xi(x) = T\tilde{\xi}(x)$ otherwise. In words, $T\xi$ is obtained by inserting a ‘flipped’ version of γ into $T\tilde{\xi}$ at site a .

To check (3.3), let us first note that if $x(\gamma^1) = x(\gamma)$, then $ST\xi$ is simply obtained by inserting the slot configurations of the sites in γ^1 into $ST\tilde{\xi}$, so that $D_k(\xi, x) = D_k(\tilde{\xi}, x)$ for $x < a$, $D_k(\xi, x + 2m) = D_k(\tilde{\xi}, x)$ for $x \geq a$ and $D_k(\xi, x) = D_k(\tilde{\xi}, a - 1)$ for $a \leq x < a + 2m$. In other words, putting γ back does not modify $D_k(\xi, x)$ in real terms. A similar relationship holds between $\mathcal{D}_k(\xi, x)$ and $\mathcal{D}_k(\tilde{\xi}, x)$. In that case, (3.3) follows from the induction hypothesis.

Now suppose that $x(\gamma^1) \neq x(\gamma)$. From Lemma 1.6 and the fact that γ^1 is the smallest soliton in $T\xi$, we deduce the only possibility to be $\gamma^1 = [a + m, a + 3m]$ with $\mathcal{H}(\gamma^1)$ preceding $\mathcal{T}(\gamma^1)$. Appealing again




Figure 3.1. We depict $T\xi$ below ξ . This example illustrates the various situations discussed in the proof of Theorem 3.1. For instance, the 5-soliton (in pink) stays put and is nested in the 6-soliton (in blue). The displacement of the 6-soliton brings five ‘new’ 5-slots to the left of the 5-soliton. (color online)

to the conservation of solitons, we see that the configuration $T\xi$ is obtained by inserting an m -soliton into $T\xi$ at the site $a + m$. It follows that $ST\xi$ is obtained by inserting the slot configurations of the sites in γ^1 into $ST\xi$ at the site $a + m$. In particular, the slot configurations of $T\xi$ and $T\xi$ coincide up to the site $a + m$. For $x \in [a + m, a + 2m)$, note that $D_k(\xi, x) = D_k(\xi, a + m - 1) + \#S_k T\xi \cap [a + m, x) - \#S_k \xi \cap [a + m, x) = D_k(\xi, a + m - 1)$, because the slot configuration in $[a + m, a + 2m)$ is unchanged. On the other hand, one can readily check $\mathcal{D}_k(\xi, x) = \mathcal{D}_k(\xi, a + m - 1)$. Finally, let us check (3.3) for $x \in [a + 2m, a + 3m)$. Writing $i = x - a - 2m + 1$ and letting A_x be the event $S\xi(x) < k$, we note that $D_k(\xi, x) - D_k(\xi, x - 1) = -\mathbb{1}\{i \leq k\} + \mathbb{1}\{A_x\}$. On the other hand, when comparing $\mathcal{D}(\xi, x)$ and $\mathcal{D}(\xi, x - 1)$, we see that if $i \leq k$, then $\mathcal{D}(\xi, x)$ loses one element, namely, $z = \mathcal{H}_i(\gamma)$, because $x = \mathcal{T}_i(\gamma^1)$; but if A_x occurs, then $x = \mathcal{H}_j(\tilde{\gamma})$ for some other soliton $\tilde{\gamma}$ and $j \leq k$, so that $x \in \mathcal{D}_k(\xi, x)$. Putting all the pieces together, we conclude that (3.3) also holds in this case, which completes its proof.

To see why (3.3) leads to the first statement in the theorem, let γ be a k -soliton to the right of $r(\xi, 0) = 0$. Let us check $D_k(\xi, x(\gamma^1)) = k$.

Firstly, suppose that $x(\gamma^1) = x(\gamma) = a$. By Lemma 1.6, this can only happen if $I(\gamma) \subseteq I(\gamma')$ for some γ' . Let $\tilde{\gamma}$ be the largest soliton among such γ' , which exists by virtue of Lemma 1.5. Because γ has to be appended to a k -slot, and because it can be checked that there is no k -slot in $[x(\tilde{\gamma}), \mathcal{H}_k(\tilde{\gamma})]$, we deduce that $\mathcal{H}_1(\tilde{\gamma}) < \dots < \mathcal{H}_k(\tilde{\gamma}) < a$. Noting that $\mathcal{T}(\tilde{\gamma}^1)$ is on the right of $\max I(\tilde{\gamma}) - 1$, we see that $\mathcal{H}_1(\tilde{\gamma}), \dots, \mathcal{H}_k(\tilde{\gamma}) \in \mathcal{D}_k(\xi, a)$. On the other hand, $\tilde{\gamma}$ is the only soliton in ξ containing γ that is appended to a record and therefore is also the only one that moves forward. If $\tilde{\gamma}$ is another soliton satisfying $x(\tilde{\gamma}) = \min I(\tilde{\gamma}) < a < \max I(\tilde{\gamma}^1)$, then we must have $I(\tilde{\gamma}) \cap I(\gamma) = \emptyset$. In particular, γ is nested in the tail of $\tilde{\gamma}^1$; that is, $a > \mathcal{T}_1(\tilde{\gamma}^1)$. Moreover, the first k -slot after $\mathcal{T}_1(\tilde{\gamma}^1)$ is $\mathcal{T}_{k+1}(\tilde{\gamma}^1)$. This implies $a \geq \mathcal{T}_{k+1}(\tilde{\gamma}^1)$. We then conclude with $D_k(\xi, a) = \#\mathcal{D}_k(\xi, a) = k$.

Secondly, suppose that $a' = x(\gamma^1) > x(\gamma)$. Let us show $\mathcal{D}_k(\xi, a') = \{\mathcal{H}_1(\gamma), \dots, \mathcal{H}_k(\gamma)\}$ in this case. Indeed, if $\widehat{\gamma} \neq \gamma$ is a soliton with $x(\widehat{\gamma}) < a' < \max I(\widehat{\gamma}^1)$, because γ is appended to a record, we must have $I(\widehat{\gamma}) \cap I(\gamma) = \emptyset$ and $\widehat{\gamma}$ has size $> k$. As previously, we deduce that γ^1 is appended to a k -slot inside the tail of $\widehat{\gamma}^1$, so that $a' \geq \mathcal{T}_{k+1}(\widehat{\gamma}^1)$. It follows that $\mathcal{H}_i(\widehat{\gamma}) \notin \mathcal{D}_k(\xi, a')$ for each $i \leq k$, and therefore $D_k(\xi, a') = k$.

Now if $r(T\xi, 0) = r(\xi, 0)$, then Proposition 1.3 and Lemma 1.5 imply that $o_k^1(\xi) = 0$ for all k . In the case that $x(\gamma^1) = x(\gamma)$, the first statement in the theorem readily follows. If, on the other hand, $x(\gamma^1) > x(\gamma)$, denote by \mathcal{T}_1 the first site in the tail of γ ; then according to Proposition 1.3 and Lemma 1.6, we must have $x(\gamma^1) = \mathcal{T}_1$. However, $[x(\gamma), \mathcal{T}_1]$ can only contain solitons of size $\leq k$ by Lemma 1.5, from which it follows that $S_k\xi \cap [x(\gamma), x(\gamma^1)] = \emptyset$. So the first statement in the theorem also holds in this case. If $r(T\xi, 0) < r(\xi, 0)$, then we need to incorporate a further shift into the indices of k -slots, which is precisely given by $o_k^1(\xi)$. If γ is to the left of $r(\xi, 0)$, the arguments are similar: it suffices to note that after each iteration, the soliton loses k k -slots on the right.

To extend to an infinite configuration $\xi \in \mathcal{W}$, let us take $x \in \mathbb{Z}$ and let $r(x)$ be the rightmost record of ξ preceding x . We note that $T\xi|_{[r(x), x]}$ only depends on $\xi|_{[r(x), x]}$ in the sense that if ξ' is a finite configuration with $\xi'(y) = \xi(y)$ for all $y \in [r(x), x]$ and $r(x)$ being a record of ξ' , then $T\xi(y) = T\xi'(y)$ for all $y \in [r(x), x]$. Because solitons are contained in excursions, we also have $S\xi(y) = S\xi'(y)$ for $y \in [r(x), x]$. Let $r'(x) = r(T\xi, -\xi(r(x))) \leq r(x)$ be the record in $T\xi$ that replaces $r(x)$. Then the same argument implies that $ST\xi(x)$ only depends on $T\xi|_{[r'(x), x]}$, which in its turn only depends on $\xi|_{[r(r'(x)), x]}$. Therefore, if we take a large enough box $[-n, n]$ containing $[r(r'(x)), x]$ as well as $r(r(T\xi, 0))$, then the restriction of ξ to this box is enough to determine the truncated k -slot configurations $(s_k(\xi, i))_{i \leq i_0}$ with $i_0 = \max\{i : s_k(\xi, i) \leq x\}$ and $(s_k(T\xi, i))_{i \leq i'_0}$ with $i'_0 = \max\{i : s_k(T\xi, i) \leq x\}$. Hence, the first statement of the theorem also holds for a general $\xi \in \mathcal{W}$.

For the third statement, suppose that $\xi, \xi' \in \mathcal{W}$ satisfy $(M_m\xi : m > k) = (M_m\xi' : m > k)$; let us show that $J_m^1\xi = J_m^1\xi'$ for all $m > k$, from which it will follow that $o_k^t(\xi) = o_k^t(\xi')$ by (3.1). We have seen that there exists some $k_0(\xi) \in \mathbb{N}$ (respectively $k_0(\xi') \in \mathbb{N}$) such that $J_m^1\xi = 0$ for all $m \geq k_0(\xi)$ (respectively $J_m^1\xi' = 0$ for all $m \geq k_0(\xi')$). Taking $k_0 = \max(k_0(\xi), k_0(\xi'))$, we see that $J_m^1\xi = J_m^1\xi' = 0$ for all $m \geq k_0$. For $k = k_0 - 1$, we immediately deduce $o_k^1(\xi) = o_k^1(\xi') = 0$ by (3.1). Moreover, we have $J_k^1\xi$ counting the number of k -solitons crossing record 0, namely, the k -solitons appended to a k -slot with a negative index at time 0 and then appended to a positive-indexed k -slot at time 1. Thus, $J_k^1\xi = \sum_{i=0}^{k-1} M_k T\xi(i) = \sum_{i=-k}^{-1} M_k \xi(i)$, by the second statement of the theorem. A similar identity holds for $J_k^1\xi'$. Therefore, if $k = k_0 - 2$, we have $J_m^1\xi = J_m^1\xi'$ for all $m > k$ and so $o_k^t(\xi) = o_k^t(\xi')$. Proceeding by downward induction, the same will be true for $k = k_0 - 2, k_0 - 3, \dots, 2, 1$. This concludes the proof of Theorem 3.1. □

4. Invariant measures

It is known that stationary Markov chains on $\{0, 1\}$ with density of 1s less than $\frac{1}{2}$ are T -invariant,¹ but in fact there are many other invariant measures for the BBS. This is due to the existence of many conservation laws intrinsic to this dynamics, in particular, the conservation of solitons studied in the previous section

In this section we prove Theorem 1.5. More precisely, we show explicitly how invariant measures can be constructed by specifying the distribution of each k th component ζ_k .

We will refer to probability measures as simply *measures*. We will also refer to measurable functions as *random elements* and refer to the pushforward of a prespecified measure by such functions as the *law* of these random elements. A measure μ on \mathcal{X} is *T-invariant* if $\mu \circ T^{-1} = \mu$.

¹The operator T applied to walks ξ is equivalent to the $2M - X$ Pitman operator. Using reversibility and Burke arguments, [HMO01] showed that product measures and stationary Markov chains with density less than $\frac{1}{2}$ are T -invariant. Extensions and a discussion of the relation between BBS and the Pitman operator can be found in [CKST18, CS19, CS20].

4.1. Construction of the measures

Our recipe to produce invariant measures uses the construction described in Subsection 2.2, and it gives a distribution $\widehat{\mu}$ of configurations ‘seen from a typical record’. The proof of Theorem 1.5 is based on properties of $\widehat{\mu}$ and how they relate to the dynamics. This relationship is given by the Palm theory, which we briefly recall now.

Let $\eta \in \{0, 1\}^{\mathbb{Z}}$. From (1.1) and (1.8), we have $\theta R\eta = R\theta\eta$. Suppose that $\eta \in \mathcal{X}$, and denote

$$r(\eta) := \inf\{x \geq 1 : x \in R\eta\}.$$

If η is random with law $\widehat{\mu}$ such that $\widehat{\mu}(r) := \int r(\eta) \widehat{\mu}(d\eta) < \infty$, its *inverse-Palm measure* $\mu = \text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu})$ is defined as follows. For every test function φ ,

$$\int \varphi(\eta) \mu(d\eta) := \int \frac{\sum_{i=0}^{r(\eta)-1} \varphi(\theta^i \eta)}{r(\eta)} \cdot \frac{r(\eta)}{\widehat{\mu}(r)} \widehat{\mu}(d\eta). \tag{4.1}$$

Informally, to sample a configuration distributed as μ one can first sample a configuration using the distribution $\widehat{\mu}$ biased by the length of the first excursion (when $0 \in R\eta$) and then choose a site uniformly from this excursion (along with the record preceding it) to place the origin.

For $\eta \in \mathcal{X}$, we define the *record-shift operator* $\widehat{\theta}$ as

$$\widehat{\theta}\eta := \theta^{r(\eta)}\eta.$$

We also define the *dynamics seen from a record* $\widehat{T} : \mathcal{X} \rightarrow \mathcal{X}$ by

$$\widehat{T}\eta := \theta^{r(T\xi^{\circ[\eta],0})}T\eta, \tag{4.2}$$

where $\xi^{\circ[\eta]}$ was defined in Subsection 2.2 as the unique lift of η with record 0 at 0. In words, \widehat{T} means apply T and recenter the configuration at the new position of record 0. The following lemma follows from standard properties of Palm measures and is proved in Section 5.

Lemma 4.3. *Let $\widehat{\mu}$ be a probability measure of $\{0, 1\}^{\mathbb{Z}}$. Suppose that $\widehat{\mu}(0 \in R\eta) = 1$, $\widehat{\mu}(r) < \infty$, $\widehat{\mu}$ is $\widehat{\theta}$ -invariant and $\widehat{\mu}(R\eta = \mathbb{Z}) = 0$. Then $\widehat{\mu}$ is supported on $\widehat{\mathcal{X}}$, and $\text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu})$ is θ -invariant, supported on \mathcal{X} and satisfies*

$$\text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu}) \circ T^{-1} = \text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu} \circ \widehat{T}^{-1}).$$

If, moreover, $\widehat{\mu}$ is $\widehat{\theta}$ -ergodic, then $\text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu})$ is θ -ergodic.

We now introduce the assumptions and notation used throughout the rest of this section.

Let $\zeta = (\zeta_k)_{k \geq 1}$ be a sequence of independent random elements of $(\mathbb{Z}_+)^{\mathbb{Z}}$ with shift-invariant distributions. Let P and E denote the underlying probability measure and expectation and suppose that the law of ζ satisfies $\sum_k kE[\zeta_k(0)] < \infty$ and $P(\sum_{k,i} \zeta_k(i) > 0) = 1$. In particular, this implies that the random field ζ a.s. satisfies (2.3) by Borel–Cantelli. Take $\xi = M^{-1}\zeta$ as the walk reconstructed from ζ according to the algorithm described in Subsection 2.2 and depicted in Figure 2.6. Let $\widehat{\mu}$ denote the resulting law of $\eta = \eta[\xi]$.

Proposition 4.4. *The measure $\widehat{\mu}$ defined above is \widehat{T} -invariant and $\widehat{\theta}$ -invariant, and it also satisfies $\widehat{\mu}(r) < \infty$ and $\widehat{\mu}(R\eta = \mathbb{Z}) = 0$. If, moreover, $(\zeta_k(i))_{i \in \mathbb{Z}}$ is i.i.d. for each k , then $\widehat{\mu}$ is also $\widehat{\theta}$ -ergodic.*

Before giving the proof, let us see how it implies Theorem 1.5.

Proof of Theorem 1.5. Let $\widehat{\mu}$ be the law of $\eta = M^{-1}\zeta$, and define

$$\mu := \text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu}).$$

By Proposition 4.4 and Lemma 4.3, μ is θ -invariant and supported on \mathcal{X} . Also,

$$\mu \circ T^{-1} = \text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu} \circ \widehat{T}^{-1}) = \text{Palm}_{\mathbb{R}}^{\mathbb{Z}}(\widehat{\mu}) = \mu,$$

so μ is also T -invariant. Moreover, under the i.i.d. assumption, the second part of Proposition 4.4 combined with the second part of Lemma 4.3 implies that μ is θ -ergodic. \square

Remark 4.5 (i.i.d. measures). A natural question is whether the product measures ν_λ can be constructed in this way. This is indeed the case, as shown in [FG20a].

Example 4.6 (Ergodicity without independent components). It is possible for law μ of η to be θ -ergodic and T -invariant while its components $M_k\eta$ are not independent under $\widehat{\mu}$. Let ζ' be the configuration $\zeta'(x) = \mathbb{1}\{x \bmod 3 = 0\}$. Let $\zeta_1 = \zeta_4$ be a configuration chosen uniformly at random in the set $\{\zeta', \theta\zeta', \theta^2\zeta'\}$; let $\zeta_k \equiv 0$ for all $k \notin \{1, 4\}$ and $\zeta = (\zeta_k)_{k \geq 1}$. The reader can check that this example satisfies the stated properties. The resulting configuration is periodic for θ and invariant for T :

$$\dots 0010111101000010000101111010000100 \dots,$$

where colors (online) represent records, 4-solitons and 1-solitons.

Example 4.7 (Independent components without ergodicity). It is also possible for ζ to be independent over k , θ -ergodic for each k but produce (by the above procedure) a configuration η whose law is not θ -ergodic. To see that, take $\zeta_5(x) \equiv 1$, ζ_1 as in the previous example, $\zeta_k \equiv 0$ for all $k \notin \{1, 5\}$ and $\zeta = (\zeta_k)_{k \geq 1}$. The resulting configurations give three classes periodic for θ and cyclic over T :

$$\begin{aligned} &\dots 0101111011000100001011110110001000 \dots \\ &\dots 0111011100100001001110111001000010 \dots \\ &\dots 0110111101000010001101111010000100 \dots \end{aligned}$$

where colors (online) represent records, 5-solitons and 1-solitons.

Remark 4.8. A measure μ is said to be θ -mixing if for all $A, B \in \mathcal{B}$, $\mu(A \cap \theta^{-n}B) \rightarrow \mu(A)\mu(B)$ as $n \rightarrow \infty$. We conjecture that, for every measure μ supported on \mathcal{X} , μ is T -invariant and θ -mixing if and only if under $\widehat{\mu}$ the components ζ_k are independent over k and each one is θ -mixing.

4.2. Invariance of the reconstructed configuration

We now prove the main part of Proposition 4.4, namely, $\widehat{\theta}$ -invariance and \widehat{T} -invariance of $\widehat{\mu}$, as well as θ -ergodicity in case of i.i.d. components. The proof of $\widehat{\mu}(\mathbf{x}) < \infty$ is given in Subsection 4.3. The condition $\widehat{\mu}(R\eta = \mathbb{Z}) = 0$ is easily checked from the construction of $M^{-1}\zeta$ and the assumption $\sum_{k,i} \zeta_k(i) > 0$ a.s. Denote by \mathcal{F}_k the sigma-field generated by $(\zeta_m)_{m > k}$.

Because one can write $\eta = M^{-1}M\eta$ and $\widehat{T}\eta = M^{-1}M\widehat{T}\eta$ by Theorem 1.3, it suffices to show that the slot decompositions $M\eta$ and $M\widehat{T}\eta$ have the same law. More precisely, it suffices to show

$$E \left(\prod_{k=1}^n \varphi_k(M_k\widehat{T}\eta) \right) = \prod_{k=1}^n E \varphi_k(\zeta_k) \tag{4.9}$$

for test functions $\varphi_1, \dots, \varphi_n$ and $n \in \mathbb{N}$.

We proceed by an induction on n . First, recall the notation from (4.2) and note that

$$\begin{aligned} &E(\varphi_k(M_k\widehat{T}\eta) \mid \mathcal{F}_k) \\ &= E(\varphi_k(\theta^{-o_k^1}(\xi^{\circ}[\eta])^{-k} M_k\eta) \mid \mathcal{F}_k) \quad (\text{by Theorem 3.1}) \\ &= E(\varphi_k(\theta^{-o_k^1}(\xi^{\circ}[\eta])^{-k} \zeta_k) \mid \mathcal{F}_k) \quad (\text{because } M_k\eta = \zeta_k) \\ &= E\varphi_k(\zeta_k), \end{aligned}$$

because ζ_k is shift-invariant and independent of $(\zeta_m)_{m>k}$, whereas $o_k^1(\xi^\circ[\eta])$ is determined by these elements. The inductive step to show (4.9) is then

$$\begin{aligned} E\left(\prod_{i=k}^n \varphi_i(M_i \widehat{T}\eta)\right) &= E\left(E\left(\prod_{i=k}^n \varphi_i(M_i \widehat{T}\eta) \middle| \mathcal{F}_k\right)\right), \\ &= E\left(\prod_{i=k+1}^n \varphi_i(M_i \widehat{T}\eta) E\left(\varphi_k(M_k \widehat{T}\eta) \middle| \mathcal{F}_k\right)\right), \\ &= E\varphi_k(\zeta_k) E\left(\prod_{i=k+1}^n \varphi_i(M_i \widehat{T}\eta)\right); \end{aligned}$$

in the second identity we have used that $M_i \widehat{T}\eta$ is determined by $(\zeta_m)_{m \geq i}$ by the last statement of Theorem 3.1. This shows that $\widehat{\mu}$ is \widehat{T} -invariant.

We now prove $\widehat{\theta}$ -invariance. Consider the transformation $M^{-1} : \zeta \mapsto \xi^*$ defined in Subsection 2.2 and denote $\eta^* = \eta[\xi^*]$ the corresponding ball configuration. Call ε_*^0 the excursion of η^* between records 0 and 1. The construction of Subsection 2.2 gives

$$\widehat{\theta}\eta^* = \theta^{r(\xi^\circ[\eta^*], 1)}\eta^* = M^{-1}(\theta^{n_k(\varepsilon_*^0)}\zeta_k : k \geq 1).$$

So it suffices to show that $(\theta^{n_k(\varepsilon_*^0)}\zeta_k)_{k \geq 1}$ has the same law as $(\zeta_k)_{k \geq 1}$. But $n_k(\varepsilon_*^0)$ is determined by $(\zeta_m)_{m>k}$ and thus is independent of ζ_k . Hence, the law of ζ_k is invariant by the random shift of $n_k(\varepsilon_*^0)$ and it is independent of $(\zeta_m)_{m>k}$. This shows that η^* and $\widehat{\theta}\eta^*$ have the same law.

Finally, under the extra assumption that $(\zeta_k(i))_{i \in \mathbb{Z}}$ is i.i.d. for each k , the reconstruction map mentioned above will produce an i.i.d. sequence of excursions separated by records. This in turn implies that the resulting configuration η is $\widehat{\theta}$ -ergodic.

4.3. Expected excursion length

We continue with the proof of Proposition 4.4 to prove that $\widehat{\mu}(\mathbf{r}) < \infty$. Let

$$\alpha_k := E[\zeta_k(0)]. \tag{4.10}$$

The proof consists of two steps: first, we show that the following system

$$w_k = 1 + \sum_{m>k} 2(m-k)w_m\alpha_m, \quad k = 0, 1, 2, \dots \tag{4.11}$$

has a unique finite nonnegative solution $w = (w_k)_{k \geq 0}$; that is, $w \in [0, \infty)^{\mathbb{N}_0}$; second, we show that the average number of k -slots per excursion in $M^{-1}\zeta$ is w_k , whence the average number of k -solitons per excursion satisfies

$$\rho_k = \alpha_k w_k. \tag{4.12}$$

In particular, this will imply that the average size of the excursions (along with the record preceding them) satisfies

$$\widehat{\mu}(\mathbf{r}) = w_0 = 1 + \sum_{m \geq 1} 2m \rho_m < \infty. \tag{4.13}$$

So we start by studying (4.11). Let $c_k := 2\sum_{m>k}(m-k)\alpha_m$. Because $\sum_k kE[\zeta_k(0)] < \infty$, we can take \bar{k} such that $\sum_{m>k-4m}\alpha_m < 1$, so $c_k < \frac{1}{2}$ for $k \geq \bar{k}$. Let $K := \{k \in \mathbb{N} : k \geq \bar{k}\} \cup \{\mathfrak{N}\}$ and consider a Markov

chain $(X_n)_{n \geq 0}$ on K with absorbing state \mathfrak{N} and transition probabilities $q_{k,m} := 2(m - k)\alpha_m \mathbb{1}\{m > k\}$; $q_{k,\mathfrak{N}} = 1 - c_k$; $q_{\mathfrak{N},\mathfrak{N}} = 1$ and $q_{k,m} = 0$ otherwise. Define the absorption time by $\tau := \inf\{n \geq 0 : X_n = \mathfrak{N}\}$. Denote by P_k the law of $(X_n)_{n \geq 0}$ starting from k and by E_k the expectation. Applying the Markov property at time 1, we see that $w_k := E_k \tau$ satisfies

$$\begin{aligned} w_k &= E_k [\tau \mathbb{1}\{X_1 = \mathfrak{N}\}] + \sum_{m>k} E_k [\tau \mathbb{1}\{X_1 = m\}] \\ &= (1 - c_k) + \sum_{m>k} (1 + w_m)q_{k,m}. \end{aligned}$$

That is, $(w_k)_{k \geq \tilde{k}}$ verifies the system (4.11) with $k = \tilde{k}, \tilde{k} + 1, \dots$. Because $c_k \geq c_{k+1}$, we have $P_k(\tau > n) \leq c_k^n$ and thus $w_k = E_k \tau \leq 1/(1 - c_k) < 2 < \infty$ for $k \geq \tilde{k}$. Because $w_{\tilde{k}} < \infty$, we can use (4.11) with $k = \tilde{k} - 1$ to define $w_{\tilde{k}-1} < \infty$, and iterating this argument we get $w_k < \infty$ for all k . This proves the existence of a finite solution to (4.11).

For uniqueness, suppose that $(\tilde{w}_k)_{k \geq 0}$ is a finite nonnegative solution to (4.11). In particular, $\sum_{m>k} \tilde{w}_m q_{k,m} = \sum_{m>k} 2(m - k)\tilde{w}_m \alpha_m < \infty$, $k \geq 0$. Moreover, the previous choice of \tilde{k} ensures that $\sum_{m>k} q_{k,m} < \infty$ for all $k \geq \tilde{k}$. It follows that for all $k \geq \tilde{k}$,

$$|w_k - \tilde{w}_k| \leq \sum_{m>k} q_{k,m} |w_m - \tilde{w}_m| \leq c_k \sup_{m>k} |w_m - \tilde{w}_m| \leq \frac{1}{2} \sup_{m>k} |w_m - \tilde{w}_m|.$$

Taking the supremum over $k \geq \tilde{k}$, we get

$$\sup_{k \geq \tilde{k}} |w_k - \tilde{w}_k| \leq \frac{1}{2} \sup_{k \geq \tilde{k}} |w_k - \tilde{w}_k|.$$

By (4.11), \tilde{w}_k and w_k are decreasing in k , so the above supremum is finite; hence, it is zero. That is, $\tilde{w}_k = E_k \tau$, $k \geq \tilde{k}$. Given $(\tilde{w}_k)_{k \geq \tilde{k}}$, (4.11) determines in a unique way the values of $(\tilde{w}_k)_{k \leq \tilde{k}}$; this completes the proof of uniqueness.

We now consider truncated approximations for the reconstruction algorithm of Subsection 2.2. Let

$$\zeta_k^{[n]}(i) := \begin{cases} \zeta_k(i), & k \leq n, \\ 0, & k > n. \end{cases}$$

Let $\varepsilon^{[n]}$ denote the first excursion (i.e., the one between records 0 and 1) of $M^{-1} \zeta^{[n]}$. Let $W_k^n \zeta = n_k(\varepsilon^{[n]})$ be the number of k -slots in $\varepsilon^{[n]}$; see (2.4). Then $W_k^n \zeta \nearrow W_k \zeta$ a.s., where $W_k \zeta := W_k^\infty \zeta$. Letting $w_k^n := E[W_k^n \zeta]$ and $w_k := E[W_k \zeta]$, by monotone convergence we have $w_k^n \nearrow w_k$ as $n \rightarrow \infty$. On the other hand, because each m -soliton contains $2(m - k)$ k -slots,

$$W_k^n \zeta = 1 + \sum_{m>k} 2(m - k) \times (\text{number of } m\text{-solitons in } \varepsilon^{[n]})$$

and thus

$$w_k^n = 1 + \sum_{m>k} 2(m - k) E(\text{number of } m\text{-solitons in } \varepsilon^{[n]}).$$

Let $\alpha_m^n := \alpha_m \mathbb{1}\{m \leq n\}$ denote the expected number of m -solitons per m -slot in $\zeta^{[n]}$. Because $W_k^n \zeta$ is a function of $(\zeta_m : m > k)$ that is independent of ζ_k , the expected number of m -solitons in $\varepsilon^{[n]}$ is $w_m^n \times \alpha_m^n$. Therefore, $(w_k^n)_{k \geq 0}$ and $(\alpha_k^n)_{k \geq 1}$ satisfy relation (4.12) and the system (4.11). Finally, because $w_k^n < 2$ for all $k \geq \tilde{k}$ and $n \in \mathbb{N}$, w_k is finite for every k and therefore (4.13) is satisfied, concluding the proof that $\tilde{\mu}(\mathbf{r}) < \infty$.

5. Palm transformations

We recall some fundamental properties on the Palm measures from Thorisson [Tho00] and study their interplay with operator T following Harris [Har71]. Although Thorisson deals with processes in the continuum, adapting the arguments to discrete space is straightforward.

For our use in Section 6, we will consider here a general situation of which $R\eta$ is a particular case. Let $Z\eta$ be a subset of \mathbb{Z} that depends on η in a translation-covariant way ($Z\theta\eta = \theta Z\eta$).

The map $\text{Palm}_{\mathbb{Z}}^Z : \mu \mapsto \tilde{\mu}$ is defined as follows. Let μ be a θ -invariant probability measure on $\{0, 1\}^{\mathbb{Z}}$ and assume that $\mu(Z\eta = \emptyset) = 0$. Define

$$\tilde{x} = \tilde{x}(\eta) := \inf\{x \geq 1 : x \in Z\eta\}$$

and $\tilde{\theta} := \theta^{\tilde{x}}$, the shift to the next element in $Z\eta$.

We define the measure $\tilde{\mu}$ as follows: for each $m \in \mathbb{N}$ and each test function φ ,

$$\int \varphi(\eta) \tilde{\mu}(d\eta) := \frac{\int \sum_{i=0}^{m-1} \varphi(\theta^i \eta) \mathbb{1}\{0 \in Z\theta^i \eta\} \mu(d\eta)}{\int \sum_{i=0}^{m-1} \mathbb{1}\{0 \in Z\theta^i \eta\} \mu(d\eta)}.$$

The above definition does not depend on m (Theorem 8.3.1 of [Tho00]). In particular, if we specialise it to the case $m = 1$, it becomes

$$\int \varphi(\eta) \tilde{\mu}(d\eta) := \frac{\int \varphi(\eta) \mathbb{1}\{0 \in Z\eta\} \mu(d\eta)}{\int \mathbb{1}\{0 \in Z\eta\} \mu(d\eta)}, \tag{5.1}$$

which means

$$\tilde{\mu} = \mu(\cdot \mid 0 \in Z\eta). \tag{5.2}$$

Theorem 8.4.1 and Formula (8.4.14) of [Tho00] then assert that the measure $\tilde{\mu}$ is $\tilde{\theta}$ -invariant. Moreover, by Formula (8.4.6) of the same book we have that the mean distance between successive points of $Z\eta$ under $\tilde{\mu}$ is the inverse density of $Z\eta$ under μ :

$$\tilde{\mu}(\tilde{x}) = \frac{1}{\mu(0 \in Z\eta)} < \infty.$$

Conversely, suppose that $\tilde{\mu}$ is a $\tilde{\theta}$ -invariant measure on $\{0, 1\}^{\mathbb{Z}}$ satisfying $\tilde{\mu}(\tilde{x}) < \infty$. Then its inverse Palm measure $\mu = \text{Palm}_{\mathbb{Z}}^Z \tilde{\mu}$ is defined as follows: for every test function φ ,

$$\int \varphi(\eta) \mu(d\eta) := \frac{\int \sum_{i=0}^{\tilde{x}(\eta)-1} \varphi(\theta^i \eta) \tilde{\mu}(d\eta)}{\int \sum_{i=0}^{\tilde{x}(\eta)-1} \mathbb{1} \tilde{\mu}(d\eta)}. \tag{5.3}$$

Moreover, μ is θ -invariant (Theorem 8.4.1 and Formula (8.4.14 °) of [Tho00]), $\mu(Z\eta = \emptyset) = 0$, and its Palm measure $\text{Palm}_{\mathbb{Z}}^Z \mu$ is given by $\tilde{\mu}$.

The above observations give the following.

Lemma 5.4. *The operations (5.2) and (5.3) define a bijection between θ -invariant measures μ on $\{0, 1\}^{\mathbb{Z}}$ with $\mu(Z\eta = \emptyset) = 0$ and $\tilde{\theta}$ -invariant measures $\tilde{\mu}$ on $\{0, 1\}^{\mathbb{Z}}$ with $\tilde{\mu}(\tilde{x}) < \infty$.*

We now analyse how the Palm transform relates to almost-sure properties.

Lemma 5.5. *Let μ be a θ -invariant measure on $\{0, 1\}^{\mathbb{Z}}$ with $\mu(Z\eta = \emptyset) = 0$ and A be a θ -invariant event. Then $\mu(A) = 0$ if and only if $\tilde{\mu}(A) = 0$.*

Proof. If $\mu(A) = 0$, then $\tilde{\mu}(A) = 0$ by (5.2). Now suppose that $\tilde{\mu}(A) = 0$. Then, again by (5.2), we have $\mu(A \cap \{0 \in Z\eta\}) = 0$. By θ -invariance of both μ and A and θ -covariance of Z , this gives $\mu(A \cap \{x \in Z\eta\}) = 0$ for every $x \in \mathbb{Z}$. Taking union over x , this gives $\mu(A) \leq \mu(Z\eta = \emptyset) = 0$. \square

As a side remark, the denominator in (5.3) is simply $\tilde{\mu}(\tilde{r})$, and this is the same formula as (4.1) with Z instead of R . We wrote it in this apparently cumbersome way to highlight the similarity with (5.1). The number 1 in the denominator of (5.3) equals the indicator that $0 \in \mathbb{Z}$ and the absence of a sum over i in (5.1) is due to the fact that the analogue to r is just the distance from 0 to the next point in \mathbb{Z} , which is 1. This also helps explain (5.7).

Lemma 5.6. *For θ -invariant μ with $\mu(Z\eta = \emptyset) = 0$, μ is θ -ergodic if and only if $\tilde{\mu}$ is $\tilde{\theta}$ -ergodic.*

Proof. Suppose that $\tilde{\mu}$ is $\tilde{\theta}$ -ergodic. Let us prove the θ -ergodicity of μ . It suffices to show that the Cesàro limits on test functions are μ -a.s. constant. Let $\eta \in \{0, 1\}^{\mathbb{Z}}$ be such that $Z\eta$ is bi-infinite. Write $k_n = \#Z\eta \cap (0, n)$. For a nonnegative bounded test function φ , we rewrite its partial average as

$$\frac{\sum_{x=0}^{n-1} \varphi(\theta^x \eta)}{n} = \frac{1}{k_n} \cdot \frac{\sum_{i=1}^{k_n-1} \sum_{j=0}^{\tilde{r}(\tilde{\theta}^i \eta)-1} \varphi(\theta^j \tilde{\theta}^i \eta) + \sum_{x=0}^{k_1-1} \varphi(\theta^x \eta) + \sum_{x=k_n}^{n-1} \varphi(\theta^x \eta)}{\sum_{i=1}^{k_n-1} \tilde{r}(\tilde{\theta}^i \eta) + k_1 + (n - k_n)}.$$

Applying the ergodic theorem [Cou16, Chapter 2] to $\tilde{\mu}$ and $\tilde{\theta}$, we find that the event A given by the set of all configurations η such that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{x=0}^{n-1} \varphi(\theta^x \eta) = \frac{\int \sum_{j=1}^{\tilde{r}(\eta)} \varphi(\theta^j \eta) \tilde{\mu}(d\eta)}{\int \tilde{r}(\eta) \tilde{\mu}(d\eta)}$$


satisfies $\tilde{\mu}(A) = 1$ (because the second and third terms of the numerator are bounded by the summands with $i = 0$ and $i = k_n$, and the same holds for the denominator). By Lemma 5.5, this proves that μ is θ -ergodic.

The implication in the other direction is proved similarly. A brief sketch is

$$\frac{\sum_{i=0}^{k-1} \varphi(\tilde{\theta}^i \eta)}{k} = \frac{\frac{1}{n_k} \sum_{x=0}^{n_k-1} \tilde{\varphi}(\tilde{\theta}^x \eta)}{\frac{k}{n_k}} \xrightarrow[k \rightarrow \infty]{\mu\text{-a.s.}} \frac{\int \tilde{\varphi} d\tilde{\mu}}{\mu(0 \in Z\eta)},$$

where $\tilde{\varphi}(\eta) := \varphi(\eta) \mathbb{1}\{0 \in Z\eta\}$ and n_k is the position of the k th element of $Z\eta$. \square

We now describe a rather useful consequence of the above theory, to be used in Section 6. It is given by the following diagram:



where μ , $\hat{\mu}$ and $\tilde{\mu}$ respectively describe the configuration seen from a typical site, a typical record or a typical element of Z . More precisely, suppose that $\mu(Z\eta = \emptyset) = \mu(R\eta = \emptyset) = 0$. Then by Lemma 5.5

the same holds for $\tilde{\mu}$ and $\widehat{\mu}$. In this case, $\tilde{\mu}$ and $\widehat{\mu}$ are related by

$$\int \varphi(\eta) \tilde{\mu}(d\eta) = \frac{\int \sum_{i=0}^{r(\eta)-1} \varphi(\theta^i \eta) \mathbb{1}\{0 \in Z\theta^i \eta\} \widehat{\mu}(d\eta)}{\int \sum_{i=0}^{r(\eta)-1} \mathbb{1}\{0 \in Z\theta^i \eta\} \widehat{\mu}(d\eta)}. \tag{5.7}$$

For the inverse transform Palm_Z^R , the same formula is valid if we swap $\tilde{\mu}, Z, \tilde{r}$ with $\widehat{\mu}, R, r$. Both (5.7) and its inverse can be proved using (5.3) and (5.1).

We finally move to studying the interplay between a Palm measure and the operator T .²

Let $\mathcal{Z} := \{\eta \in \mathcal{X} : Z\eta \text{ is bi-infinite}\}$. Suppose for each $\eta \in \mathcal{Z}$ that there is a bijection Ψ_η between $Z\eta$ and $ZT\eta$ that depends on η in a translation covariant way; that is, $\Psi_{\theta\eta}(\theta x) = \theta\Psi_\eta(x)$. Intuitively, Ψ is just an honest way to follow elements of $Z\eta$ after applying the operator T . Let

$$\tilde{\mathcal{Z}} := \{\eta \in \mathcal{Z} : 0 \in Z\eta\}.$$

For $\eta \in \tilde{\mathcal{Z}}$, we define

$$\tilde{T}\eta = \theta^{\Psi_\eta(0)} T\eta.$$

We now show the following.

Lemma 5.8. *For θ -invariant μ on \mathcal{Z} , we have the identity*

$$\text{Palm}_{\tilde{\mathcal{Z}}}^{\mathcal{Z}}(\mu \circ T^{-1}) = (\text{Palm}_{\tilde{\mathcal{Z}}}^{\mathcal{Z}} \mu) \circ \tilde{T}^{-1}.$$

In particular, μ is T -invariant if and only if $\text{Palm}_{\tilde{\mathcal{Z}}}^{\mathcal{Z}} \mu$ is \tilde{T} -invariant.

Proof. We follow the classical arguments in [Har71, PS73]. For a test function φ , we have

$$\begin{aligned} \int \mathbb{1}\{0 \in ZT\eta\} \varphi(T\eta) \mu(d\eta) &= \int \sum_x \mathbb{1}\{x \in Z\eta, \Psi_\eta(x) = 0\} \varphi(T\eta) \mu(d\eta) \\ &= \sum_x \int \mathbb{1}\{0 \in Z\theta^x \eta, \Psi_{\theta^x \eta}(0) = -x\} \varphi(T\eta) \mu(d\eta) \\ &= \sum_x \int \mathbb{1}\{0 \in Z\eta, \Psi_\eta(0) = -x\} \varphi(T\theta^{-x} \eta) \mu(d\eta) \\ &= \int \mathbb{1}\{0 \in Z\eta\} \sum_x \mathbb{1}\{\Psi_\eta(0) = x\} \varphi(T\theta^x \eta) \mu(d\eta) \\ &= \int \mathbb{1}\{0 \in Z\eta\} \varphi(\tilde{T}\eta) \mu(d\eta). \end{aligned}$$

The third identity holds by the translation invariance of μ . Taking $\varphi \equiv 1$, we get $(\mu \circ T^{-1})(\tilde{\mathcal{Z}}) = \mu(\tilde{\mathcal{Z}})$. Hence,

$$\begin{aligned} \text{Palm}_{\tilde{\mathcal{Z}}}^{\mathcal{Z}}(\mu \circ T^{-1}) \varphi &= \frac{1}{(\mu \circ T^{-1})(\tilde{\mathcal{Z}})} \int \mathbb{1}\{0 \in ZT\eta\} \varphi(T\eta) \mu(d\eta) \\ &= \frac{1}{\mu(\tilde{\mathcal{Z}})} \int \mathbb{1}\{0 \in Z\eta\} \varphi(\tilde{T}\eta) \mu(d\eta) \\ &= ((\text{Palm}_{\tilde{\mathcal{Z}}}^{\mathcal{Z}} \mu) \circ \tilde{T}^{-1}) \varphi. \end{aligned} \quad \square$$

²For concreteness, we consider the BBS operator T , but we only use the fact that it is a θ -covariant operator and its domain \mathcal{X} is θ -invariant.

Proof of Lemma 4.3. Define a bijection ψ_η between $R\eta$ and $RT\eta$ as follows. For $x \in R\eta$, let $\psi_\eta(x) := r(T\xi, j)$, where $\xi = \xi[\eta]$ and $x = r(\xi, j)$. Note that this definition does not depend on the lift $\xi[\eta]$. Note also that $\psi_{\theta\eta}(\theta x) = \theta\psi_\eta(x)$. Finally, note that (4.2) can be written as

$$\widehat{T}\eta := \theta^{\psi_\eta(0)}T\eta.$$

All of the statements in the lemma follow from the previous lemmas by taking $Z\eta = R\eta$ and $\Psi = \psi$, except that μ is supported on \mathcal{X} and $\widehat{\mu}$ is supported on $\widehat{\mathcal{X}}$. By Lemma 5.5, it is enough to show the latter.

By the ergodic theorem, the limit

$$\lambda(\eta) = \lim_{y \rightarrow \infty} \frac{1}{y} \sum_{x=-y}^{-1} \eta(x) = \lim_{y \rightarrow \infty} \frac{1}{y} \sum_{x=1}^y \eta(x)$$

exists for μ -almost everywhere (a.e.) η , hence for $\widehat{\mu}$ -a.e. η , and we have to show that $0 < \lambda(\eta) < \frac{1}{2}$ for $\widehat{\mu}$ -a.e. η .

By the ergodic decomposition theorem [Cou16, Chapter 14], we can assume that $\widehat{\mu}$ is $\widehat{\theta}$ -ergodic. In the above limit, we can take a subsequence $y_k \rightarrow \infty$ with elements from $R\eta$, giving

$$\lambda(\eta) = \lim_{k \rightarrow \infty} \frac{\sum_{x=1}^{y_k} \eta(x)}{y_k} = \lim_{k \rightarrow \infty} \frac{\sum_{j=0}^{k-1} \frac{1}{2}(\mathbf{r}(\widehat{\theta}^j \eta) - 1)}{\sum_{j=0}^{k-1} \mathbf{r}(\widehat{\theta}^j \eta)} = \frac{\frac{1}{2}(\widehat{\mu}(\mathbf{r}) - 1)}{\widehat{\mu}(\mathbf{r})} < \frac{1}{2}.$$

By assumption, $\widehat{\mu}(R\eta = \mathbb{Z}) = 0$; thus, $\widehat{\mu}(\mathbf{r} = 1) < 1$ and hence $\widehat{\mu}(\mathbf{r}) > 1$ and the above limit is also positive, concluding the proof. □

6. Asymptotic speed of solitons

In this section we prove Theorems 1.1 and 1.2 from a combination of simpler statements. By looking at the dynamics as seen from a k -soliton, we show that the soliton speed v_k exists and equals the expected length of the jump of a typical k -soliton in one step.

In Subsection 6.1 we show that the soliton speed v_k that appears in (1.9) is μ -a.s. well defined and is given by an explicit formula (6.1). We also show that it is given another formula (6.3) and that it is finite. Analysing the interaction between solitons of different sizes, in Subsection 6.2 we show that the soliton speeds $(v_k)_k$ satisfy (1.10). For completeness, in Subsection 6.3 we show that v_k is positive and increasing in k without assuming that $\bar{\rho}_k > 0$. In Subsection 6.4 we analyse the formula (6.3) using the description of μ from Subsection 4.1 to show that the soliton speeds are given by (1.12). Finally, in Subsection 6.5 we briefly mention the results about vertical speeds.

6.1. Existence of speeds via Palm measure and ergodicity

Here we show that $\lim_t \frac{1}{t}x(\gamma^t)$ exists μ -a.s. and we give an explicit formula for it. Recall the definition of $x(\gamma)$ given before the statement of Theorem 1.1. Let $\Gamma_k^\circ \eta := \{x(\gamma) : \gamma \in \Gamma_k \eta\}$ denote the set of leftmost sites of k -solitons of η . If $\rho_k = 0$, then Γ_k is empty and (1.9) holds for all $\gamma \in \Gamma_k(\eta)$ by vacuity, for any value of v_k .

From now on, we assume $\rho_k > 0$, which implies $\mu(0 \in \Gamma_k^\circ \eta) > 0$. Because $\Gamma_k^\circ \theta \eta = \theta \Gamma_k^\circ \eta$, the Palm construction described in Section 5 applies.

For $\gamma \in \Gamma_k \eta$ and $z = x(\gamma)$, we define $\Delta_\eta^k(z) := x(\gamma^1) - x(\gamma)$, the size of the jump of k -soliton γ after one iteration of T . For $z \notin \Gamma_k^\circ \eta$ we set $\Delta_\eta^k(z) = 0$. With this notation, the displacement of a tagged k -soliton after $t + 1$ iterations of T can be decomposed as

$$x(\gamma^{t+1}) - x(\gamma) = \Delta_\eta^k(x(\gamma)) + \Delta_{T\eta}^k(x(\gamma^1)) + \dots + \Delta_{T^t \eta}^k(x(\gamma^t)).$$

We want to divide both sides by t and use the ergodic theorem. This will require a couple of subtle observations. The first step is to consider the system as seen from a typical k -soliton.

Let $\widehat{\mathcal{X}}^k$ be the set of configurations in \mathcal{X} such that $\Gamma_k^\circ \eta$ is doubly infinite and contains 0. Let $\widehat{\mu}_k := \text{Palm}_{\mathbb{Z}}^{\Gamma_k^\circ}(\mu)$ be the Palm measure of μ with respect to $\Gamma_k^\circ \eta \subseteq \mathbb{Z}$; that is, $\widehat{\mu}_k = \mu(\cdot | 0 \in \Gamma_k^\circ \eta)$. For $\eta \in \widehat{\mathcal{X}}^k$, let $r^k(\eta) := \inf\{x \geq 1 : x \in \Gamma_k^\circ \eta\}$ and define $\widehat{\theta}_k : \widehat{\mathcal{X}}^k \rightarrow \widehat{\mathcal{X}}^k$ as the ‘shift to the next k -soliton’ given by $\widehat{\theta}_k \eta := \theta^{r^k(\eta)} \eta$. Also, for a k -soliton γ such that $x(\gamma) = 0$, let $\widehat{T}_k \eta := \theta^{x(\gamma^1)} T \eta$ denote the dynamics as seen from a tagged k -soliton.

From Lemma 5.8, $\widehat{\mu}_k$ is \widehat{T}_k -invariant. Now for $\eta \in \widehat{\mathcal{X}}^k$ and γ with $x(\gamma) = 0$, the above decomposition becomes

$$x(\gamma^{t+1}) = \Delta_\eta^k(0) + \Delta_{\widehat{T}_k \eta}^k(0) + \Delta_{\widehat{T}_k^2 \eta}^k(0) + \dots + \Delta_{\widehat{T}_k^t \eta}^k(0).$$

By the ergodic theorem for $(\widehat{T}_k^t \eta)_{t \in \mathbb{Z}}$, $\lim_{t \rightarrow \infty} \frac{1}{t} x(\gamma^t)$ exists $\widehat{\mu}_k$ -a.s., and on average it equals

$$v_k := \int \Delta_\eta^k(0) \widehat{\mu}_k(d\eta) \tag{6.1}$$

(to avoid integrability issues, note that by Lemma 1.6 we have $\Delta_\eta^k(0) \geq 0$).

It remains to show that this limit is in fact nonrandom.

Consider the field

$$\tilde{v}_k(\eta, z) := \begin{cases} \lim_{t \rightarrow \infty} \frac{x(\gamma^t)}{t}, & \text{if } z = x(\gamma) \text{ for some } \gamma \in \Gamma_k(\eta), \\ 0, & \text{otherwise.} \end{cases}$$

By Proposition 1.3 and Lemma 1.5, if $\gamma, \tilde{\gamma}$ are two k -solitons with $x(\gamma) \leq x(\tilde{\gamma})$, then $x(\gamma^t) \leq x(\tilde{\gamma}^t)$ for all t . Hence, we have

$$\tilde{v}_k(\eta, x) \leq \tilde{v}_k(\eta, y) \text{ for all } x \leq y \text{ in } \Gamma_k^\circ \eta. \tag{6.2}$$

On the other hand, by Lemma 5.6 the measure $\widehat{\mu}_k$ is $\widehat{\theta}_k$ -ergodic. So, under $\widehat{\mu}_k$ the sequence $(\tilde{v}_k(\widehat{\theta}_k^j \eta, 0))_{j \in \mathbb{Z}} \in [0, +\infty]^{\mathbb{Z}}$ is a.s. nondecreasing by (6.2) and its law is θ -ergodic, which implies that it is $\widehat{\mu}_k$ -a.s. equal to some constant $v_k \geq 0$. So $\tilde{v}_k(\eta, x) = v_k$ for all $x \in \Gamma_k^\circ(\eta)$, for $\widehat{\mu}_k$ -a.e. η . By Lemma 5.5, it also holds for μ -a.e. η .

We conclude with a short proof that $v_k < \infty$. From (5.7) with Γ_k° in the role of Z , we get

$$v_k = \frac{\int \sum_{i=0}^{r(\eta)-1} \mathbb{1}\{\theta^i \eta \in \widehat{\mathcal{X}}^k\} \Delta_{\theta^i \eta}^k(0) \widehat{\mu}(d\eta)}{\int \sum_{i=0}^{r(\eta)-1} \mathbb{1}\{\theta^i \eta \in \widehat{\mathcal{X}}^k\} \widehat{\mu}(d\eta)} = \frac{1}{\rho_k} \int \sum_{y=0}^{r(\eta)-1} \Delta_\eta^k(y) \widehat{\mu}(d\eta). \tag{6.3}$$

Denote by $h(\gamma)$ the leftmost site of the tail of γ and by $t(\gamma)$ the leftmost site of its head. Then Proposition 1.3 implies that $\Delta_\eta^k(x) \leq |h(\gamma) - t(\gamma)|$. Combined with the fact that intervals spanned by different k -solitons do not overlap (Lemma 1.5), this yields $\sum_{y=1}^{r(\eta)} \Delta_\eta^k(y) \leq r(\eta)$, from which we get $v_k \leq \frac{\widehat{\mu}(r)}{\rho_k} = \frac{w_0}{\rho_k} < \infty$ by (4.13).

6.2. Equation for speeds from soliton interactions

We now prove that the soliton speeds satisfy (1.10). It is enough to consider k such that $\rho_k > 0$ and otherwise take (1.10) as the definition of v_k . We start by proving the following identity on the displacement of solitons.

Proposition 6.4. *We have*

$$x(\gamma^t) - x(\gamma) = k t - 2k \sum_{m>k} \frac{1}{2} M_t^\gamma(m) + \sum_{m<k} 2m N_t^\gamma(m), \tag{6.5}$$

where

$$M_t^\gamma(m) = \#\{(\tilde{\gamma}, s) : \tilde{\gamma} \in \Gamma_m \eta, 0 \leq s < t, I(\gamma^s) \subseteq I(\tilde{\gamma}^s) \text{ and } x(\tilde{\gamma}^{s+1}) \neq x(\tilde{\gamma}^s)\},$$

$$N_t^\gamma(m) = \#\{(\tilde{\gamma}, s) : \tilde{\gamma} \in \Gamma_m \eta, 0 \leq s < t, I(\tilde{\gamma}^s) \subseteq I(\gamma^s) \text{ and } I(\tilde{\gamma}^{s+1}) \cap I(\gamma^{s+1}) = \emptyset\}.$$

Proof. Before proving it, let us first explain the intuition underpinning (6.5). From Lemma 1.6, we know that as long as there exists some $\tilde{\gamma}$ satisfying $I(\gamma) \subseteq I(\tilde{\gamma})$, γ is prevented from moving. We will see that it takes two time steps for $\tilde{\gamma}$ to overtake γ : the first step occurs when γ is nested in the right half of $\tilde{\gamma}$, after which it will be nested in the left half of $\tilde{\gamma}$; the second step occurs when γ is nested in the left half of $\tilde{\gamma}$, after which it will no longer be nested inside $\tilde{\gamma}$. Because $\tilde{\gamma}$ can be inside an even larger soliton, during which it stays frozen, these two times are not necessarily consecutive. On the other hand, when γ overtakes a smaller soliton of size ℓ say, it causes γ to move 2ℓ units more than it normally would have. Put formally, we have the following identity:

$$x(\gamma^1) - x(\gamma) = k - k \cdot \mathbb{1}\{\exists \tilde{\gamma} : I(\gamma) \subseteq I(\tilde{\gamma}) \text{ and } x(\tilde{\gamma}^1) \neq x(\tilde{\gamma})\} + \sum_{\ell < k} 2\ell N_0^\gamma(\ell). \tag{6.6}$$

The proof of this identity is elementary and is postponed to Section 7. Proceeding with the proof of (6.5), we note that if $\tilde{\gamma}$ exists in (6.6), then it has to be unique: solitons $\tilde{\gamma}$ satisfying $I(\tilde{\gamma}) \supseteq I(\gamma)$ are ordered by their sizes and only the largest one will move forward. Iterating the above t times then yields (6.5). \square

Now take $\tilde{\gamma}$ to be the leftmost m -soliton in $[1, \infty)$ and γ the rightmost k -soliton in $(-\infty, 0]$. We have shown previously that $v_m = \lim_{t \rightarrow \infty} \frac{1}{t} x(\tilde{\gamma}^t)$ and $v_k = \lim_{t \rightarrow \infty} \frac{1}{t} x(\gamma^t)$, μ -a.s. Combined with Lemma 1.7, this yields

$$v_k \leq v_m \quad \text{if } k < m \text{ and } \rho_k, \rho_m > 0.$$

From now on, we will work under $\widehat{\mu}_k$ and always take γ to be the soliton with $x(\gamma) = 0$. To lighten the notation, we will drop the superscript γ from M_t and N_t . Recall that $\bar{\rho}_m = \frac{\rho_m}{w_0}$ is the mean number of m -solitons per unit space.

Let us assume for the moment that $\bar{\rho}_m = 0$ for all large m .

By (1.9) and (6.5), to get (1.10) it suffices to show the following. For $m > k$ and $\ell < k$:

$$\begin{aligned} \frac{1}{2t} M_t(m) &\rightarrow \bar{\rho}_m (v_m - v_k), \\ \frac{1}{t} N_t(\ell) &\rightarrow \bar{\rho}_\ell (v_k - v_\ell), \end{aligned} \quad \text{in probability.} \tag{6.7}$$

Proof of (6.7). Let us denote $\mathbf{M}_t(m) = \{(\tilde{\gamma}, s) : \tilde{\gamma} \in \Gamma_m \eta, 0 \leq s < t, I(\gamma^s) \subseteq I(\tilde{\gamma}^s) \text{ and } x(\tilde{\gamma}^{s+1}) \neq x(\tilde{\gamma}^s)\}$, the set of solitons that overtake γ , so that $M_t(m) = \#\mathbf{M}_t(m)$. Let γ_- be the leftmost m -soliton satisfying $x(\gamma_-) > x(\gamma)$ and let γ_+ be the rightmost m -soliton satisfying $x(\gamma_+) < x(\gamma)$. Let us consider the m -solitons $\tilde{\gamma}$ that are found between γ^t and γ^t_- at time t and denote by \mathbf{L}_t^- their set, namely, $\mathbf{L}_t^- := \{\tilde{\gamma} \in \Gamma_m \eta : x(\gamma^t) < x(\tilde{\gamma}^t) < x(\gamma^t_-)\}$. Note that if $\tilde{\gamma} \in \mathbf{L}_t^-$, then necessarily $I(\gamma^t) \cap I(\tilde{\gamma}^t) = \emptyset$. We also define $\mathbf{L}_t^+ = \{\tilde{\gamma} \in \Gamma_m \eta : I(\tilde{\gamma}^t) \cap [x(\gamma^t), x(\gamma^t_+)] \neq \emptyset\}$, which includes in particular the m -soliton $\tilde{\gamma}$ satisfying $I(\gamma^t) \subseteq I(\tilde{\gamma}^t)$ (if such a soliton exists). Roughly speaking, $\#\mathbf{L}_t^-$ counts those m -solitons that have completed the 2-step overtaking of γ between times 0 and $t - 1$, so it will give an undercount of $M_t(m)$, whereas $\#\mathbf{L}_t^+$ will be an overcount. More precisely, let us show that $2\#\mathbf{L}_t^- \leq M_t(m) \leq 2\#\mathbf{L}_t^+$. Indeed, if $(\tilde{\gamma}, s) \in \mathbf{M}_t(m)$, then Lemma 1.7 implies that necessarily $x(\tilde{\gamma}) < x(\gamma)$. It follows that $x(\tilde{\gamma}) \leq x(\gamma_+)$, and then $x(\tilde{\gamma}^t) \leq x(\gamma^t_+)$. On the other hand, because $I(\gamma^s) \subseteq I(\tilde{\gamma}^s)$, there are only two

possibilities at $s + 1$: either $x(\tilde{\gamma}^{s+1}) > x(\gamma^{s+1})$ ($\tilde{\gamma}$ overtakes γ), which implies $x(\tilde{\gamma}^t) > x(\gamma^t)$ by Lemma 1.7, or $I(\gamma^{s+1}) \subseteq I(\tilde{\gamma}^{s+1})$. In the latter case, if $s + 1 < t$, we then repeat the arguments used for s ; otherwise, we have $I(\gamma^t) \subseteq I(\tilde{\gamma}^t)$. All cases considered, we conclude that $\tilde{\gamma} \in \mathbf{L}_t^+$. Moreover, for such a $\tilde{\gamma}$, there are at most two times s such that $(\tilde{\gamma}, s) \in \mathbf{M}_t(m)$: the first time s' when $I(\gamma^{s'}) \subseteq I(\tilde{\gamma}^{s'})$ and $\tilde{\gamma}$ is the largest soliton containing γ (i.e., first step in overtaking) and the first time s'' after s' when $\tilde{\gamma}$ moves forward (i.e., second step in overtaking). And for all $u \geq s'' + 1$, we have $x(\tilde{\gamma}^u) > x(\gamma^u)$ as implied by Lemmas 1.6 and 1.7, so that γ and $\tilde{\gamma}$ will never intersect again. Because s' and s'' do not necessarily belong to $[0, t)$, we only have $M_t(m) \leq 2\#\mathbf{L}_t^+$. For the other inequality, we note that if $\tilde{\gamma} \in \mathbf{L}_t^-$, then $x(\tilde{\gamma}) < x(\gamma_-)$, which implies actually $x(\tilde{\gamma}) \leq x(\gamma_+) < x(\gamma)$. Because $x(\tilde{\gamma}^t) > x(\gamma^t)$, it follows that both s' and s'' described previously belong to $[0, t)$. So the claimed inequalities follow.

Write whp to denote ‘with high $\widehat{\mu}_k$ -probability’ and note that high μ implies high $\widehat{\mu}_k$. By (1.9), for every $\varepsilon > 0$, $x(\gamma^t)$ is in $v_k t \pm \varepsilon t$ and $x(\gamma_-^t)$ is in $v_m t \pm \varepsilon t$ whp. Using θ -invariance of μ , by the ergodic theorem applied to counting m -solitons on a large interval, the number of m -solitons of $T^t \eta$ located in $[v_k t \pm \varepsilon t, v_m t \pm \varepsilon t]$ is within $\bar{\rho}_k(v_m - v_k)t \pm 3\varepsilon t$ whp. Therefore, $\#\mathbf{L}_t^-$ is in $\bar{\rho}_m(v_m - v_k)t \pm 3\varepsilon t$ whp. Because \mathbf{L}_t^+ and \mathbf{L}_t^- can differ by at most two m -solitons, this concludes the proof that $\frac{1}{2t}M_t(m) \rightarrow \bar{\rho}_m(v_m - v_k)$ in probability.

The proof for $N_t(\ell)$ is similar but simpler. The only difference is that for each $\tilde{\gamma} \in \Gamma_\ell \eta$, there is at most one time s such that $(\tilde{\gamma}, s)$ is counted in the tally $N_t(\ell)$. □

To complete the proof of (1.10), it remains to drop the assumption that $\bar{\rho}_m = 0$ for all large m . The above proof contains all of the argument, except that combining (6.5) and (6.7) requires a limit and an infinite sum to commute. To prove the general case, we take expectation in (6.5) with $t = 1$ and combine it with (6.1):

$$\begin{aligned} v_k &= \int x(\gamma^1) \widehat{\mu}_k(d\eta) = \int \left[k - \sum_{m>k} k M_1(m) + \sum_{m<k} 2m N_1(m) \right] \widehat{\mu}_k(d\eta) \\ &= k - k \sum_{m>k} \int M_1(m) \widehat{\mu}_k(d\eta) + \sum_{m<k} 2m \int N_1(m) \widehat{\mu}_k(d\eta). \end{aligned}$$

It remains to show $\int N_1(m) \widehat{\mu}_k(d\eta) = \bar{\rho}_m(v_k - v_m)$ and $\frac{1}{2} \int M_1(m) \widehat{\mu}_k(d\eta) = \bar{\rho}_m(v_m - v_k)$. Similar to the argument in Subsection 6.1, we decompose

$$N_{t+1}(m) = N_1^\eta(m) + N_{\widehat{T}_k^\eta}^1(m) + N_{\widehat{T}_k^2 \eta}^1(m) + \dots + N_{\widehat{T}_k^\eta}^1(m)$$

with a slight abuse of notation. By \widehat{T}_k -invariance, $\frac{1}{t}N_t^\eta(m)$ converges $\widehat{\mu}_k$ -a.s. to a random variable (i.e., a measurable function of η) whose average is $\int N_1^\eta(m) \widehat{\mu}_k(d\eta)$. On the other hand, by (6.7) this variable is constant and equal to $\bar{\rho}_m|v_k - v_m|$. The same argument works for $M_1(m)$. This concludes the proof of (1.10) and also of Theorem 1.1.

6.3. Soliton speeds are positive and increasing

Here we show that v_k is positive and increasing in k without assuming that $\bar{\rho}_k > 0$. Both follow from combining (1.10) with the following bound on the total flow of solitons crossing the origin:

$$2\bar{\rho} \cdot v := 2 \sum_m \bar{\rho}_m v_m < 1. \tag{6.8}$$

The idea is that $\rho_m v_m$ is the average flow of solitons through the origin and each one takes two time steps. Before sketching the proof, we show that soliton speeds are positive and increasing.

Splitting and regrouping the sums in (1.10) gives

$$v_k = k - \sum_m 2(k \wedge m)\bar{\rho}_m v_m + v_k \sum_m 2(k \wedge m)\bar{\rho}_m = av_k + b,$$

where $a \leq 2\lambda < 1$ and $b \geq k - 2k\bar{\rho} \cdot v > 0$. Hence, $v_k > 0$.

Now, writing $d_k := v_{k+1} - v_k > 0$, by subtracting the above equation from itself we get

$$d_k = 1 + d_k \sum_m 2(k \wedge m)\bar{\rho}_m + \sum_{m \geq k+1} 2\bar{\rho}_m(v_{k+1} - v_m) = ad_k + b + c,$$

where $a \leq 2\lambda < 1$, $b \geq 1 - 2\bar{\rho} \cdot v > 0$ and $c = \sum_{m > k} 2\bar{\rho}_m v_{k+1} \geq 0$. Therefore, $d_k > 0$.

Proof of (6.8). The proof bears much similarity to that of (6.7), so we only sketch the arguments here. Denote by $Q_t(k) = \#\{(\gamma, s) : \gamma \in \Gamma_k \eta, 0 \leq s < t, 0 \in I(\gamma^s) \text{ and } x(\gamma^{s+1}) \neq x(\gamma^s)\}$ the number of occurrences that a k -soliton is the largest one straddling 0 between times 0 and t . Note that each k -soliton that crosses the origin in the time interval $[0, t]$ contributes twice to $Q_t(k)$, once when the origin is in its tail and once in its head.

To find the growth rate of $Q_t(k)$ as $t \rightarrow \infty$, let us denote by γ_+ the rightmost k -soliton γ' satisfying $x(\gamma') \leq 0$. The counts $Q_t^-(k) = 2\#\{\gamma \in \Gamma_k \eta : 0 < x(\gamma^t) < x(\gamma_+^t)\}$ and $Q_t^+(k) = 2\#\{\gamma \in \Gamma_k \eta : I(\gamma^t) \cap [0, x(\gamma_+^t)] \neq \emptyset\}$ satisfy $Q_t^-(k) - 3 \leq Q_t(k) \leq Q_t^+(k)$. Arguments identical to those in (6.7) show that $\frac{1}{t}Q_t(k) \rightarrow 2\bar{\rho}_k v_k$ in probability for each $k \geq 1$.

Applying the ergodic theorem for T , convergence of $\frac{1}{t}Q_t(k)$ says that $2\bar{\rho}_k v_k$ equals the probability that the largest soliton γ such that $0 \in I(\gamma)$ belongs to $\Gamma_k \eta$. For each $\eta \in \mathcal{X}$, either $0 \in R\eta$ or there is a unique such k , so

$$\mu(0 \in R\eta) + \sum_k 2\bar{\rho}_k v_k = 1,$$

concluding the proof. □

6.4. Recursion formulas for independent components

We now prove Theorem 1.2. Let $\zeta = M\eta$. We are assuming that the field $(\zeta_k(i))_{i \in \mathbb{Z}}$ is i.i.d. over i for each k and independent over k . So let us proceed the other way around. We let P denote the law of ζ and E the corresponding expectation. In this notation, $\hat{\mu}$ is the law of $\eta = M^{-1}\zeta$.

Note that (4.11) and (4.12) give the first two equations in (1.12). The third equation can be taken as the definition of s_k , and it is a simple recursive definition once one has ρ , w and α . Combining these with (6.3), to get the last equation in (1.12) we need to show that

$$E \sum_{y=0}^{r(\eta)-1} \Delta_\eta^k(y) = \alpha_k \cdot s_k. \tag{6.9}$$

We now use the assumption that $\eta = M^{-1}\zeta$, where M^{-1} denotes de reconstruction map of Subsection 2.2. First, note that

$$\sum_{y=0}^{r(\eta)-1} \Delta_\eta^k(y) = \sum_\gamma x(\gamma_1) - x(\gamma),$$

where the sum is over all k -solitons γ located between record 0 and record 1. Moreover, k -solitons appended to k -slots belonging to larger solitons will stay put, just switching zeros for ones, and only the k -solitons that are appended directly to the 0th k -slot at $x = 0$ will actually jump (Lemma 1.6).

Furthermore, by Proposition 1.3 and Lemma 1.6, the size of their jump equals the distance between their leftmost 1 and their leftmost 0; that is, the distance between the tip of their head and the tip of their tail.

Now the number of k -solitons appended to the 0th k -slot is exactly $\zeta_k(0)$, which on average equals α_k by (4.10). So to conclude the proof of (6.9), it is enough to observe that s_k given by the recursion relation in (1.12) in fact gives the expected distance between the tips of the head and tail of a typical k -soliton appended to a record. By independence of ζ_k over k , being appended to a record is irrelevant, and we can consider a typical k -soliton instead.

We claim that s_k equals the a.s. empirical average of the distance between the tip of the head and the tip of the tail among all k -solitons and also that $2s_k$ equals the a.s. empirical average of the size of $I(\gamma)$ among all k -solitons γ . We made the previous statement stronger so we can prove it by induction. Remember that the interval $I(\gamma)$ consists of sites occupied by γ together with the smaller solitons appended to its slots. For $k = 1$ we have $s_1 = 1$, consistent with the fact that a 1-soliton is always given by the strings $a = 10$ or $\tilde{a} = 01$ with nothing appended inside it. For $k = 2$, note that each 2-soliton (including smaller solitons appended to its slots) is of the form $b = 11\tilde{a}^*00a^*$ or $\tilde{b} = 00a^*11\tilde{a}^*$ where a^* stands for $\zeta_1(i)$ copies of a and \tilde{a}^* stands for $\zeta_1(j)$ copies of \tilde{a} for some i, j that are determined by $(\zeta_m)_{m \geq 2}$. So the average size of $11\tilde{a}^*$ and that of $00a^*$ both equal $2 + 2\alpha_1$. For $k = 3$, note that each 3-soliton is of the form $c = 11\tilde{a}^*1\tilde{a}^*\tilde{b}^*00a^*0a^*b^*$ or $\tilde{c} = 00a^*0a^*b^*11\tilde{a}^*1\tilde{a}^*\tilde{b}^*$ where b^* stands for $\zeta_2(i)$ independent copies of b , etc. So the average size of each half of a 3-soliton equals $s_3 = 3 + 2s_2\alpha_2 + 4s_1\alpha_1$. The induction step is clear, which concludes the proof of Theorem 1.2.

6.5. Vertical speed

A previous version of this work [arXiv:1806.02798v3] focussed on measuring vertical speed of solitons, before the authors realised that it was in fact possible to study the horizontal speeds.

We now briefly mention the results about vertical speed. The analysis is very similar to what was done in Subsection 6.2, and tedious details will be omitted.

For $\beta \in R\eta$, we define $\beta^t := r(T^t\xi, j)$, where $\beta = r(\xi, j)$. Note that this definition does not depend on the lift $\xi[\eta]$. Recalling (3.2), we define the displacement of a k -bearer $\pi \in S_k\eta$ measured in terms of records by

$$y_k^t(\eta, \pi) = \#\{\beta \in R\eta : \pi < \beta \text{ and } \pi^{k,t} \geq \beta^t\}, \quad \pi \in S_k\eta.$$

In case there is a k -soliton $\gamma \in \Gamma_k\eta$ appended to the k -slot π in η , the tagged k -soliton γ^t will appear appended to the k -slot $\pi^{k,t}$ in $T^t\eta$, so y_k^t also measures the displacement of tagged k -solitons, but it is well defined even when there are no k -solitons. The limit (6.13) gives a physical meaning to the soliton speeds v_k when $\rho_k = 0$ in Theorem 1.1.

Theorem 6.1. *Let μ be a measure on \mathcal{X} such that under $\hat{\mu}$ each k th component $M_k\eta$ is i.i.d. and they are independent over k . There exists a nondecreasing deterministic sequence $h = (h_k)_{k \geq 1}$ such that μ -a.s. on η , for all $k \in \mathbb{N}$ and $\pi \in S_k\eta$,*

$$\lim_{t \rightarrow \infty} \frac{y_k^t(\eta, \pi)}{t} = h_k \in [k, \infty]. \tag{6.10}$$

Assuming $\sum_k k^2 \rho_k < \infty$, the vector $(h_k)_{k \geq 1}$ is the unique finite solution of the linear system

$$h_k = k + \sum_{m > k} 2(m - k)(h_m - h_k)\rho_m, \quad k \geq 1. \tag{6.11}$$

The asymptotic speed of tagged records is given by

$$\lim_{t \rightarrow \infty} -\frac{\beta^t}{t} = v_0 := \sum_{m \geq 1} 2m\rho_m h_m, \quad \text{for all } \beta \in R\eta, \mu\text{-a.s.} \tag{6.12}$$

The asymptotic speed v_k of k -bearers is also given by

$$\lim_{t \rightarrow \infty} \frac{\pi^{k,t}}{t} = v_k = h_k w_0 - v_0, \quad \text{for all } \pi \in S_k\eta, \mu\text{-a.s.}, \tag{6.13}$$

From the two last equations, the speed of tagged records is also given by

$$v_0 = \sum_{m \geq 1} 2m\rho_m v_m. \tag{6.14}$$

Furthermore, the vertical speed of any walk representation $\xi = \xi[\eta]$ is given by

$$\lim_{t \rightarrow \infty} -\frac{T^t \xi(0)}{t} = h_0 := \frac{v_0}{w_0}. \tag{6.15}$$

Let us outline the argument for the first four equations.

The proof of (6.10) is similar to that of (1.9) and only uses θ -ergodicity of η . Each time an m -soliton overtakes a k -bearer, this causes the position of the k -bearer measured in records to be incremented by an extra factor of $2(m - k)$. On the other hand, the position of a k -bearer measured in records is not affected by overtaking smaller solitons. These two facts explain the origin of (6.11), but the proof uses a truncation argument that deletes all large components, and the i.i.d. assumption is to ensure that the resulting configuration is still θ -ergodic.

The speeds h_k are either all finite or all infinite; they are finite if and only if $\sum_m m^2 \rho_m < \infty$. Deleting large solitons, one gets from (6.11) that $0 \leq h_m - h_k \leq m - k$. Plugging this back into (6.11) and using $\sum_m m \rho_m < \infty$, one eventually gets $h_{k+j} - h_k \geq \delta j$ for some $\delta > 0$ and all large k . This gives $\delta k \leq h_k \leq k + 2 \sum_m m^2 \rho_m$ uniformly with respect to the deletion threshold, which yields stated equivalence. Also, under this condition, one can show existence and uniqueness of the solution to (6.11) using truncation, as we did for (4.11).

Each time a tagged m -soliton crosses a tagged record from left to right, it causes the record to move $2m$ boxes left. On the other hand, by mass conservation the number of such crossings by time t equals $\rho_m y_m^t$, which is about $\rho_m h_m t$ by (6.10). Summing over m we get (6.12). As a side remark, the vertical speed h_0 is given by the expected vertical jump at the origin $h_0 = \int (\xi[\eta](0) - T\xi[\eta](0))\mu(d\eta)$, and intuitively this is related to $\sum_k k^2 \rho_k$ because the excursion contained the origin is size-biased.

Finally, by (6.10), the k -bearer $\pi = 0 \in S_k \xi$ will typically have crossed about $h_k t$ records by time t , so it will be between two tagged records with initial index about $h_k t$. By ergodicity, the initial position of these records is about $w_0 h_k t$, so by (6.12) their position at time t will be about $w_0 h_k t - v_0 t$. Dividing by t and taking a limit, one gets (6.13).

From (6.13) we have $h_m = \frac{v_0 + v_m}{w_0}$, and substituting into (6.12) and using (4.13) we get (6.14). To prove (6.15), we note that after t iterations of T , record i will be at $x = o(t)$ if $r(\xi, i) = v_0 t + o(t)$, which implies that $T^t \xi(0) = -i + o(t)$. On the other hand, $r(\xi, i) = w_0 i + o(i)$, whence $T^t \xi(0) = -h_0 t + o(t)$, concluding the proof.

7. Postponed proofs

Proof of Proposition 1.3. Let us prove for finite η first. The proof is by induction on number of balls contained in η . Identifying 0 with ‘ \ominus ’ and 1 with ‘ \oplus ’, consider the following data stream version of the TS-Algorithm.

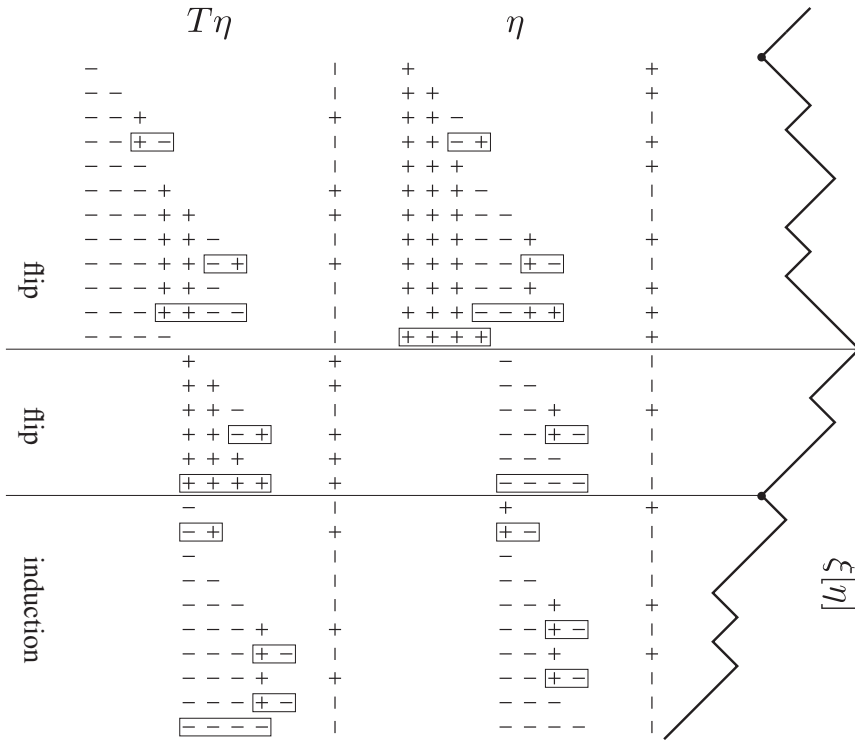


Figure 7.1. Example showing conservation of solitons by splitting space in three parts: rising, falling and remainder. After applying T , the configurations on the rising and falling parts are flipped, the smaller solitons are conserved and flipped and the biggest soliton moves forward and will have its tail in the remainder part. Applying T to the remainder part conserves solitons by induction.

Also, because $\xi(x) \geq \xi(y)$ for all $x \in I_1 \cup I_2$, by definition of T we have that η and $T\eta$ are complements of each other on these two blocks. So by the previous observations, the streaming algorithm applied to η and to $T\eta$ on I_1 will produce exactly the opposite suffixes at every step. The same is true for I_2 . The only difference is that now the \ominus^k produced after processing $T\eta$ on I_1 is incorporated into the infinite prefix \ominus^∞ , and it will not annihilate the \oplus^k obtained after processing $T\eta$ on I_2 . Hence, while processing $T\eta$ on $I_1 \cup I_2$, the same solitons will be generated, with \oplus replaced by \ominus ; that is, with the head occupying the former position of the tail, except for this last k -soliton.

Finally, the \oplus^k obtained after processing $T\eta$ on $I_1 \cup I_2$ will not increase its length while processing $T\eta$ on I_3 , because processing $T\eta$ on I_3 always produces \ominus -suffixes. So this run \oplus^k is preserved until the first time when the processing of $T\eta$ on I_3 produces a \ominus^k , and they both annihilate. This eventually occurs because $T\eta$ has infinitely many records to the right. So again the head of the corresponding k -soliton will take the position previously occupied by the tail of a k -soliton. Moreover, when it occurs, it annihilates \ominus s that were not going to be annihilated while processing $(T\eta)|_{I_3}$ because they would have been simply absorbed by the prefix \ominus^∞ . Hence, the presence of this \oplus^k does not change how the algorithm processes $T\eta$ on I_3 , either before or after such annihilation occurs. To conclude, note that $\eta|_{I_3}$ contains fewer balls than η so we can assume by induction that the tails of the solitons of $\eta|_{I_3}$ will become the heads of the solitons of $T\eta|_{I_3}$, proving the proposition for the case of a finite configuration η .

We finally consider general $\eta \in \mathcal{X}$. Let A be a set of k sites. Let y_2, y_3 be records for $T\eta$ to the left and right of A , respectively. Let $y_1 < y_2$ and $y_4 > y_3$ be records for η . Let η' denote the restricted configuration, given by $\eta'(x) := \eta(x)\mathbb{1}\{x \in [y_1, y_4]\}$. Because solitons are always contained in the interval between two consecutive records, if some $\gamma \in \Gamma_k \eta$ intersects A then it is contained in $[y_1, y_4]$. Because $\eta' \leq \eta$, and x being a record for η is a nonincreasing property of η , y_1 and y_4 are also records

for η' . Hence, the soliton configuration $\Gamma_k\eta$ restricted to $[y_1, y_4]$ coincides with $\Gamma_k\eta'$. Now notice that $T\eta' = T\eta$ on $[y_1, y_4]$ and $T\eta' = 0$ on $(-\infty, y_1]$. In particular, $T\eta' = T\eta$ on $[y_2, y_3]$, $T\eta' \leq T\eta$ on $(-\infty, y_2]$, and thus y_2, y_3 are also records for $T\eta'$. Hence, by the same argument as above, if some $\gamma \in \Gamma_k T\eta'$ intersects A , then it is contained in $[y_2, y_3]$; moreover, $\Gamma_k T\eta$ restricted to $[y_2, y_3]$ coincides with $\Gamma_k T\eta'$ restricted to $[y_2, y_3]$. Because η' is a finite configuration, by the previous case this concludes the proof. \square

Proof of Lemma 1.5. Let us first take a configuration with a finite number of balls and let γ be a soliton. The second statements follow from the definition of $I(\gamma)$. For the rest, note that while implementing the Takahashi–Satsuma algorithm, when γ is identified, it appears as a set of consecutive sites (after the previous steps that remove solitons lodged inside $I(\gamma)$). Moreover, the rules of the algorithm imply that solitons intersecting $I(\gamma)$ that are removed from the previous steps must have a strictly smaller size than γ and are completely contained in $I(\gamma)$. A formal proof can be done by an induction on the number of solitons. These properties can be extended to infinite configurations similar to the proof of Proposition 1.3, because the algorithm applies to each excursion of the configuration. \square

Proof of Lemma 1.6. Assume in the first place that there is a finite number of balls and we proceed by an induction on the number of solitons. The basic case when there is only one soliton is clear. Suppose that the statement holds true for a system with up to n solitons and suppose that η is a configuration with $n + 1$ solitons. Let k be the minimal size of its solitons and let γ be the leftmost k -soliton. Denote by $a = x(\gamma)$. Then Lemma 1.5 implies that $\mathcal{H}(\gamma) \cup \mathcal{T}(\gamma) = [a, a + 2k)$.

Let us first show $x(\gamma^1) \geq x(\gamma) = a$. If $x(\gamma^1) < a$, combined with the fact that $\mathcal{H}(\gamma^1) = \mathcal{T}(\gamma)$, this would imply $\mathcal{H}(\gamma^1) \cup \mathcal{T}(\gamma^1) = [a - k, a + k)$, because γ^1 has minimal size in $T\eta$. Moreover, in this case we would have $\mathcal{T}(\gamma^1) = \{a - k, \dots, a - 1\}$ and $\mathcal{T}(\gamma) = \{a, \dots, a + k - 1\}$. If $a - 1$ is a record of η , then a will also be a record of η because $\eta(a) = 0$, which is absurd. If $a - 1 \notin R\eta$ but $a - 2 \in R\eta$, we readily check $a + 1 \in R\eta$ in this case, which is also impossible unless $k = 1$. Proceeding in this way, we deduce that all of the sites of $\mathcal{T}(\gamma^1)$ must belong to the head of some soliton γ' , which must have size $> k$ as a result of our choice of γ . Note the site $a - k - 1$ also belongs to the head of γ' ; otherwise, the TS algorithm would produce a k -soliton from the sites $[a - k, a + k]$. Because none of the sites in $[a - k - 1, a + k]$ are records, they are flipped in $T\eta$, so that we have a run of at least $k + 1$ 0s preceding a . But this contradicts the rules of the TS algorithm to have a k -soliton on $[a - k, a + k]$. This means it is impossible to have $x(\gamma^1) < a$.

If $a - 1 \in R\eta$, then we must have $\mathcal{H}(\gamma) = [a, a + k)$ and $\mathcal{T}(\gamma) = [a + k, a + 2k)$; otherwise, we would have $k + 1$ 0s followed by k 1s, contradicting the TS algorithm. Then in $T\eta$, the sites in $[a, a + 2k)$ get flipped and $a - 1$ still has value 0. In particular, we cannot have γ^1 on the same sites as γ . But $x(\gamma^1) \geq x(\gamma)$ and $\mathcal{H}(\gamma^1) = \mathcal{T}(\gamma)$. So the only possibility is $\mathcal{T}(\gamma^1) = [a + 2k, a + 3k)$.

If instead $a - 1 \notin R\eta$, then it also gets flipped in $T\eta$. Note that we must have $\eta(a - 1), \eta(a)$ having different values; otherwise, the TS algorithm would not produce a soliton starting from a . In that case, it is straightforward to check that γ^1 is the soliton on the sites $[a, a + 2k)$ and then $\mathcal{T}(\gamma^1) = \mathcal{H}(\gamma)$.

So far we have shown $x(\gamma^1) \geq x(\gamma)$ and that $\mathcal{T}(\gamma^1) = \mathcal{H}(\gamma)$ if and only if $a - 1$, the k -slot to which γ is appended is not a record, so it must belong to some soliton γ' with size $> k$. Moreover, in this case, we have $I(\gamma) \cap I(\gamma') \neq \emptyset$, which yields $I(\gamma) \subseteq I(\gamma')$ by Lemma 1.5. Applying the induction hypothesis to the configuration η' obtained from η by removing γ , we see that the statement also holds for the n solitons of η' . Together with the previous arguments for γ , this implies that the statement holds for all of the $n + 1$ solitons of η , because if we have $I(\gamma_1) \subseteq I(\gamma_2)$, then inserting smaller solitons will not affect this. To extend to the infinite configuration, we follow the same strategy as employed in the proof of Theorem 3.1, noting that $T\eta(x)$ depends only on $\eta|_{[r(x), x]}$, with $r(x)$ being the rightmost record preceding x ; we omit the details here. \square

Proof of Lemma 1.7. By Lemma 1.5 we have $\max I(\gamma) < x(\tilde{\gamma})$. By Proposition 1.3, $x(\gamma^1) \leq \max I(\gamma)$. By Lemma 1.6, $x(\tilde{\gamma}) \leq x(\tilde{\gamma}^1)$. Combining these we have $x(\gamma^1) < x(\tilde{\gamma}^1)$, and by induction on t we get $x(\gamma^t) < x(\tilde{\gamma}^t)$. \square

Proof of (6.6). By Lemma 1.6, we have $x(\gamma^1) = x(\gamma)$ if and only if there exists some $\hat{\gamma}$ satisfying $I(\gamma) \subseteq I(\hat{\gamma})$. Moreover, Lemma 1.5 implies that if there is another soliton γ' with $I(\gamma) \subseteq I(\gamma')$, then $\hat{\gamma}$ and γ' are of different sizes and the smaller one is nested inside the larger one. In that case, we take $\tilde{\gamma}$ to be the soliton with maximal size satisfying $I(\tilde{\gamma}) \supseteq I(\gamma)$. Note that $\tilde{\gamma}$ is necessarily appended to a record (if not, we can find an even larger soliton that contains γ), and therefore $x(\tilde{\gamma}^1) \neq x(\tilde{\gamma})$.

It remains to show that if $x(\gamma^1) > x(\gamma)$, then their difference is given by $\sum_{\ell < k} 2\ell N'_0(\ell)$. According to Lemma 1.6, we must have $\mathcal{H}(\gamma)$ to the left of $\mathcal{T}(\gamma)$ in this case. It then follows from Proposition 1.3 that $x(\gamma^1)$ is the leftmost site of $\mathcal{T}(\gamma)$, denoted as $t(\gamma)$. To conclude, it suffices to note that $t(\gamma) - x(\gamma) = k + \sum_{\gamma'} 2 \times \text{size of } \gamma'$, where the sum is over the solitons γ' lodged inside the left side of γ . One can readily check that this leads to the desired identity. \square

Proof of Proposition 1.11. Let $c_k := \sum_{m \neq k} 2(m \wedge k) \bar{\rho}_m$. Because $\ell \bar{\rho}_\ell \rightarrow 0$, by dominated convergence we have $c := \sup_k c_k = \lim_k c_k = \sum_\ell 2\ell \bar{\rho}_\ell < \frac{1}{2}$ using our assumption. Define

$$q_{k,m} := \begin{cases} \frac{1}{c} 2(m \wedge k) \bar{\rho}_m & \text{if } m \neq k \\ \frac{1}{c} (c - c_k) & \text{if } m = k \end{cases} \quad \text{and} \quad Q := (q_{k,m})_{k,m}.$$

Because $\sum_m 2m \bar{\rho}_m < \infty$, we can split the sums and rewrite (1.10) as

$$v_k = k + c v_k - c (Qv)_k = \frac{k - c (Qv)_k}{1 - c}, \quad v_k \in [0, \infty). \tag{7.1}$$

Note that $(Qv)_k$ is finite for solutions v to (7.1).

Let v and \tilde{v} be two solutions to (1.10), and denote $r_k := v_k - \tilde{v}_k$. Then

$$s_k := |r_k| = a |(Qr)_k| \leq a (Qs)_k,$$

where $a := c/(1 - c) < 1$. By induction, for all $n \in \mathbb{N}$,

$$s_k \leq a^n (Q^n s)_k. \tag{7.2}$$

The measure $\pi = (\pi_k)_{k \geq 1}$ given by

$$\pi_k := \frac{\bar{\rho}_k}{\sum_\ell \bar{\rho}_\ell}$$

satisfies $\pi Q = \pi$. Because $\frac{1}{1-c} < 2$ and $q_{k,m} \geq 0$, solutions to (7.1) satisfy $v_k \leq 2k$ and

$$\pi v = \sum_k v_k \bar{\rho}_k (\sum_\ell \bar{\rho}_\ell)^{-1} \leq 2 \sum_k k \bar{\rho}_k (\sum_\ell \bar{\rho}_\ell)^{-1} < \infty.$$

Likewise, $\pi \tilde{v} < \infty$ and $\pi s \leq \pi v + \pi \tilde{v} < \infty$. Using (7.2) and Q -invariance of π ,

$$\pi_k s_k \leq \pi s \leq a^n \pi Q^n s = a^n \pi s,$$

for every $n \in \mathbb{N}$. Hence, if $\bar{\rho}_k > 0$, then $\pi_k > 0$ and $|r_k| = 0$, implying $v_k = \tilde{v}_k$. When $\bar{\rho}_k = 0$ we have v_k is a function of $(v_m : \bar{\rho}_m > 0)$, implying uniqueness also in this case. \square

Acknowledgements. We thank Roberto Fernández and Davide Gabrielli for very helpful discussions. We thank the referees for a very careful reading and suggestions that greatly improved the article.

Financial support: PAF was partially supported by Argentinian grants ANPCyT (PICT 2015-3583, 2018-02842) and UBACyT 20020160100155BA.

Conflict of Interest: None.

References

- [BDS83] C. Boldrighini, R. L. Dobrushin and Y. M. Sukhov, ‘One-dimensional hard rod caricature of hydrodynamics’, *J. Stat. Phys.* **31** (1983), 577–616.
- [CBS19] X. Cao, V. B. Bulchandani and H. Spohn, ‘The GGE averaged currents of the classical Toda chain’, *J. Phys. A* **52**(49) (2019), 495003.
- [Cou16] Y. Coudène, *Ergodic Theory and Dynamical Systems*, (Springer, London, 2016).
- [CKST18] D. A. Croydon, T. Kato, M. Sasada and S. Tsujimoto, ‘Dynamics of the box-ball system with random initial conditions via Pitman’s transformation’, Preprint 2018, arXiv:1806.02147.
- [CS19] D. A. Croydon and M. Sasada, ‘Invariant measures for the box-ball system based on stationary Markov chains and periodic Gibbs measures’, *J. Math. Phys.* **60** (2019), 083301.
- [CS20] D. A. Croydon and M. Sasada, ‘Discrete integrable systems and Pitman’s transformation’, Preprint, 2020, arXiv:2007.06206.
- [CS21] D. A. Croydon and M. Sasada, ‘Generalized hydrodynamic limit for the box-ball system’, *Comm. Math. Phys.* **383** (2021), 427–463.
- [DS17] B. Doyon and H. Spohn, ‘Dynamics of hard rods with initial domain wall state’, *J. Stat. Mech. Theory Exp.* **2017** (2017) 073210.
- [DYC18] B. Doyon, T. Yoshimura and J.-S. Caux, ‘Soliton gases and generalized hydrodynamics’, *Phys. Rev. Lett.* **120** (2018), 045301.
- [FG20a] P. A. Ferrari and D. Gabrielli, ‘BBS invariant measures with independent soliton components’, *Electron. J. Probab.* **25** (2020), 26.
- [FG20b] P. A. Ferrari and D. Gabrielli, ‘Box-ball system: soliton and tree decomposition of excursions’, in *XIII Symposium on Probability and Stochastic Processes* (Basel, Switzerland, 2020), 107–152.
- [HMO01] B. Hambly, J. B. Martin and N. O’Connell, ‘Pitman’s $2M - X$ theorem for skip-free random walks with Markovian increments’, *Electron. Comm. Probab.* **6** (2001), 73–77.
- [Har71] T. E. Harris, ‘Random measures and motions of point processes’, *Z. Wahrscheinlichkeitstheorie verw. Geb.* **18** (1971), 85–115.
- [IKO04] R. Inoue, A. Kuniba and M. Okado, ‘A quantization of box-ball systems’, *Rev. Math. Phys.* **16** (2004), 1227–1258.
- [IKT12] R. Inoue, A. Kuniba and T. Takagi, ‘Integrable structure of box-ball systems: crystal, Bethe ansatz, ultradiscretization and tropical geometry’, *J. Phys. A* **45** (2012), 073001.
- [KTZ17] T. Kato, S. Tsujimoto and A. Zuk, ‘Spectral analysis of transition operators, automata groups and translation in BBS’, *Comm. Math. Phys.* **350** (2017), 205–229.
- [KLO18] A. Kuniba, H. Lyu and M. Okado, ‘Randomized box–ball systems, limit shape of rigged configurations and thermodynamic Bethe ansatz’, *Nucl. Phys. B* **937** (2018), 240–271.
- [KMP20] A. Kuniba, G. Misguich and V. Pasquier, ‘Generalized hydrodynamics in box-ball system’, *J. Phys. A* **53** (2020), 404001.
- [LPS21] T. Lam, P. Pylyavsky and R. Sakamoto, ‘Rigged configurations and cylindric loop Schur functions’, *Ann. Inst. Henri Poincaré Comb. Phys. Interact.* Preprint, 2021, arXiv:1410.4455.
- [LLP17] L. Levine, H. Lyu and J. Pike, ‘Phase transition in a random soliton cellular automaton’, Preprint, 2017, arXiv:1706.05621.
- [MIT06] J. Mada, M. Idzumi and T. Tokihiro, ‘The exact correspondence between conserved quantities of a periodic box-ball system and string solutions of the Bethe ansatz equations’, *J. Math. Phys.* **47** (2006), 053507.
- [Pit75] J. W. Pitman, ‘One-dimensional Brownian motion and the three-dimensional Bessel process’, *Adv. Appl. Probab.* **7** (1975), 511–526.
- [PS73] S. C. Port and C. J. Stone, ‘Infinite particle systems’, *Trans. Amer. Math. Soc.* **178** (1973), 307–340.
- [Sak14a] R. Sakamoto, ‘Rigged configurations and Kashiwara operators’, *Symmetry Integr. Geom.* **10** (2014), 028.
- [Sak14b] R. Sakamoto, ‘Ultradiscrete soliton systems and combinatorial representation theory’, *RIMS Kokyuroku* **1913** (2014), 141–158.
- [Spo12] H. Spohn, *Large Scale Dynamics of Interacting Particles* (Springer Science & Business Media, Berlin, 2012).
- [TM97] D. Takahashi and J. Matsukidaira, ‘Box and ball system with a carrier and ultradiscrete modified KdV equation’, *J. Phys. A* **30** (1997), L733–L739.
- [TS90] D. Takahashi and J. Satsuma, ‘A soliton cellular automaton’, *J. Phys. Soc. Japan* **59** (1990), 3514–3519.
- [Tho00] H. Thorisson, ‘Coupling, stationarity, and regeneration’, in *Probability and its Applications* (Springer-Verlag, New York, 2000).
- [TTMS96] T. Tokihiro, D. Takahashi, J. Matsukidaira and J. Satsuma, ‘From soliton equations to integrable cellular automata through a limiting procedure’, *Phys. Rev. Lett.* **76** (1996), 3247.
- [TTS96] M. Torii, D. Takahashi and J. Satsuma, ‘Combinatorial representation of invariants of a soliton cellular automaton’, *Physica D* **92** (1996), 209–220.
- [Zuk20] A. Zuk, ‘From partial differential equations to groups’, in *Analysis and Geometry on Graphs and Manifolds* (Cambridge University Press, Cambridge, 2020), 368–381.