Soliton decomposition of the Box-Ball System

Pablo A. Ferrari, Chi Nguyen, Leonardo T. Rolla, Minmin Wang

December 7, 2018

Abstract

The Box-Ball System was introduced by Takahashi and Satsuma as a discrete counterpart of the KdV equation. Both systems exhibit solitons whose shape and speed are conserved after collision with other solitons. Conservation of solitons suggests that this dynamics has many spatially-ergodic invariant measures besides the i.i.d. distribution. Meanwhile, solitons of different sizes interact through a momentary change of speeds during collision, which cumulatively affects their asymptotic speeds, suggesting that the speeds are determined by such interaction. In order to understand general invariant measures and soliton interactions, we introduce a decomposition of configurations through slots, reducing the dynamics to a simple hierarchical translation of different components. Using this property we obtain an explicit recipe to construct a rich family of invariant measures. Finally, we obtain explicit equations for the soliton speeds in terms of spacial density of solitons.

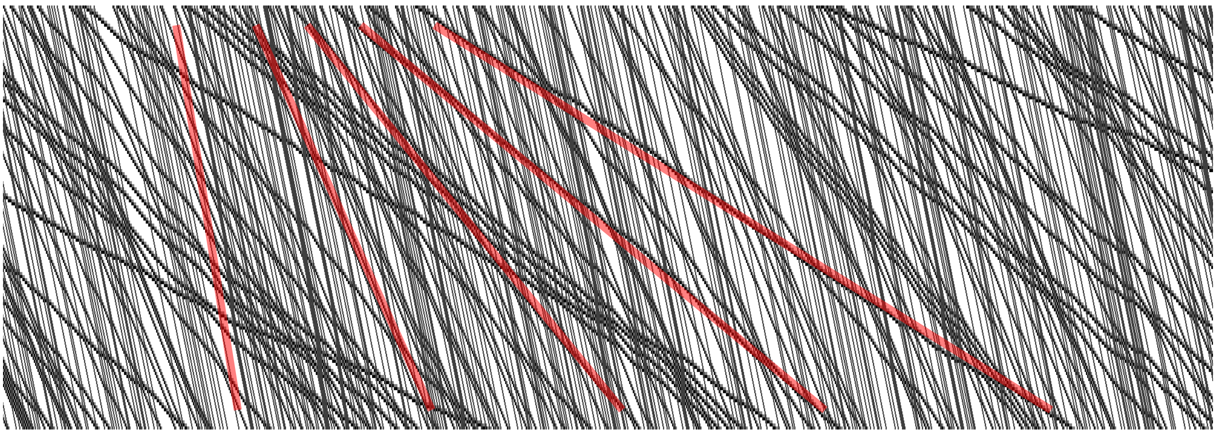

BBS dynamics for i.i.d. initial configuration with density 0.25. Straight red lines are deterministic and computed using Theorem 1.2. (high resolution, color online)

1 Introduction

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 $\eta(x) := 0$. Consider first configurations with a finite number of balls and let an empty carrier start to the left of the leftmost ball and visit the boxes one after another. When visiting box x, the carrier picks a ball if there is any and if the x is empty and the carrier has at least one ball, he deposits the ball in the box. 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 1 0 1 1 0 0 0 0 1 1 0 1 0 0 0 0 0    η
 0 0 1 0 1 2 1 0 0 0 1 2 1 2 1 0 0 0     carrier load
0 0 0 1 0 0 1 1 0 0 0 0 1 0 1 1 0 0 0    Tη
```

This dynamics, called *Box-Ball System* (BBS), was introduced in [TS90] as a discrete system showing solitons, a phenomenon also present in the KdV equation. For the relation between BBS and KdV see [TTMS96, TM97]. For further physical motivation of the BBS model, see [KTZ17, LLP17, TTS96]. See also [CKST18, IKO04, IKT12, LPS14, MIT06, Sak14a, Sak14b] for some other developments.

The transformation T can be defined for configurations with infinitely many balls if the density of balls is well-defined and less than $\frac{1}{2}$. The product measure with any such density is invariant under T. But the family of invariant probability measures is far richer than that, since this transformation T has many conservation properties. In particular, a configuration can have *solitons* of all sizes, these solitons are conserved by the dynamics and there is a distance measure between solitons of the same size which is also conserved.

Takahashi and Satsuma proposed an algorithm to identify solitons in a finite ball configuration and argued that the solitons identified at time 0 can be tracked at successive iterations of T. An isolated *k-soliton* γ consists of k successive occupied boxes followed by k successive empty boxes. Evolving this configuration by t iterations of T, the new configuration will have a k-soliton γ^t which is a translation of γ by kt. The striking property of the BBS is that, although solitons can collide which momentarily changes their shapes and introduces shifts, such collisions neither create nor destroy solitons, see Proposition 1.4.

Given a soliton size $k \geqslant 1$ and a ball configuration η, we identify a subset of boxes called *k-slots*. To each k-slot, one can append any number of k-solitons. For each ball configuration, we can describe the number of k-solitons appended to each k-slot and call *k-component* the resulting vector. The components are defined hierarchically starting from the bigger solitons, and ball configurations can be reconstructed from the components. We show that under the BBS evolution the k-component is rigid, conserving the number of k-slots between two successive k-solitons, for all k. More precisely, the k-component is shifted by a quantity depending only on the m-soliton configuration for $m > k$, see Theorem 2.1.

Our main result about invariant measures is the following. Given a family of shift-invariant probability measures on $\mathbb{N}_0^{\mathbb{Z}}$ indexed by k and whose densities decay fast enough in k, we can construct a T-invariant probability measure μ whose components are independent and distributed according to such family, see Theorem 1.3. We conjecture that this decomposition characterizes 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.

We finally study the asymptotic speed of k-solitons. The position of each k-soliton divided by time converges to a deterministic asymptotic speed v_k, and $(v_k)_{k\geqslant 1}$ is the unique solution of an explicit system of linear equations, see Theorems 1.1 and 1.2.

In the sequel we give precise definitions and state some of the main results.

Our space will be the set of configurations defined by

$$\mathcal{X}_\lambda := \Big\{\eta \in \{0,1\}^{\mathbb{Z}} : \lim_{y\to\infty} \frac{1}{y} \sum_{x=-y}^{0} \eta(x) = \lim_{y\to\infty} \frac{1}{y} \sum_{x=0}^{y} \eta(x) = \lambda\Big\} \quad \text{and} \quad \mathcal{X} := \bigcup_{0<\lambda<\frac{1}{2}} \mathcal{X}_\lambda.$$

For $\eta \in \mathcal{X}$ we define the set of *records* by

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

Sites in between records form finite *excursions*. The operator T is defined by

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

Lemma 1.3. *Let $0 \leqslant \lambda < \frac{1}{2}$ and $\eta \in \mathcal{X}_\lambda$. Then $T\eta \in \mathcal{X}_\lambda$.*

The above lemma is a direct consequence of (2.1), see §2.

1.1 Identifying solitons and computing speeds

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

The *k-solitons* are identified in [TS90] by the following algorithm.

Start with a doubly infinite *word*, so that each letter in the word is 0 or 1 and remembers which box x it corresponds to in the ball configuration η
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 run preceding it
  Remove these $2k$ letters from the *word*
end

Notice that a k-soliton consists of k zeros followed by k ones or vice-versa, and letters which do not belong to any soliton are all zero and correspond to the records of η, see Fig. 1.1. For general $\eta \in \mathcal{X}$ we apply the above algorithm to each excursion.

Let γ be a k-soliton. We define the *head* and *tail* of γ as follows: the head $\mathcal{H}(\gamma)$ is the set of the positions of k ones in γ and the tail $\mathcal{T}(\gamma)$ is formed by the positions of the k zeros. Let $\Gamma_k\eta$ be the set of k-solitons of a ball configuration $\eta \in \mathcal{X}$. The following is proved in §5.

Proposition 1.4. *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 $(\gamma^t)_{t\geqslant 0}$ the trajectory satisfying $\gamma^0 = \gamma$ and

$$\mathcal{H}(\gamma^{t+1}) = \mathcal{T}(\gamma^t). \tag{1.5}$$

Let μ be a shift-ergodic (when we say ergodic it also implies invariant) measure on $\mathcal{X}$. Denote by ρ_k the mean number of k-solitons per excursion, by $w_0 = 1 + \sum_k 2k\rho_k$ the mean size of excursions, and by $\bar\rho_k = \frac{\rho_k}{w_0}$ the mean number of k-solitons per site (precise definitions in §3.3). Denote by $x(\gamma)$ the leftmost site of γ. We now state the main result concerning speeds.

Theorem 1.1. *Let μ be a T-invariant and shift-ergodic measure on $\mathcal{X}_\lambda$. Then there exists deterministic $v = (v_k)_k$ such that, μ-a.s., for all $\gamma \in \Gamma_k\eta$,*

$$\lim_{t\to\infty} \frac{x(\gamma^t)}{t} = v_k. \tag{1.6}$$

The speeds $(v_k)_k$ are finite and satisfy the system

$$v_k = k + \sum_{m<k} 2m\bar\rho_m(v_k - v_m) - \sum_{m>k} 2k\bar\rho_m(v_m - v_k). \tag{1.7}$$

For the next result we need a stronger assumption, described further below.

Theorem 1.2. *Let μ have independent i.i.d. soliton components. Then the speeds $(v_k)_k$ in* (1.6) *are also given by the unique solution to*

$$w_k = 1 + \sum_{m>k} 2(m-k)\rho_m, \quad \rho_k = \alpha_k w_k, \quad s_k = k + \sum_{m<k} 2(k-m)s_m\alpha_m, \quad v_k = \frac{s_k}{w_k}, \tag{1.8}$$

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

So the above theorems describe the speeds explicitly. In practice, the speeds v_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

.111100~~10~~111010000110110000.10.	.111100**10**111010000110110000.10.
.1111~~0011~~1010000110110000.10.	.1111**00**10**11**1010000110110000.10.
.11111~~01~~0000110110000.10.	.11110010111**01**0000110110000.10.
.11111000011~~01~~10000.10.	.1111001011101000011**01**10000.10.
.111110000~~111000~~0.10.	.11110010111010000**11**01**1000**0.10.
.~~1111100000~~.10.	.**1111**001011**1**01**0000**11011000**0**.10.

Figure 1.1: Applying the 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 it ignores the other 1-soliton. (color online)

finally v. For i.i.d. initial conditions, one can find α_k explicitly in terms of the density λ by computing partition functions [FG18], substitute the equation for ρ into that for w, then compute s and v. Using this and the above theorem, we have produced the simulations shown in Figs. 1.2 and 1.3 as well as that of the first page.

System (1.7) 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 be delayed by 2 time steps if $m > k$. The term $\bar{\rho}_m|v_k - v_m|$ gives the frequency of such encounters as seen from a k-soliton.

In system (1.8), w_k is the density of k-slots per excursion (see below), 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 §4 for details).

1.2 Slots and invariant measures

Recall that configurations $\eta \in \mathcal{X}$ are decomposed into solitons and records, and that a k-soliton has a head and a tail, each one consisting of k (possibly non-consecutive) sites. We say that the j-th site of the head or tail of an m-soliton has a k-slot for all $k < j$. Roughly speaking, k-slots are the places where k-solitons can be inserted without interfering with the rest of the configuration in terms of the Takahashi-Satsuma algorithm. We also say that a record has k-slots for every k. Indeed, a soliton of any given size can be inserted between two records.

Assume that $0 \in R\eta$. Enumerate the k-slots in increasing order and so that the 0-*th* k-*slot* be 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 it is contained between $s_k(\eta, i)$ and $s_k(\eta, i+1)$. Define the k-*component* of η as the configuration $M_k\eta$ of k-solitons appended to the k-slots,

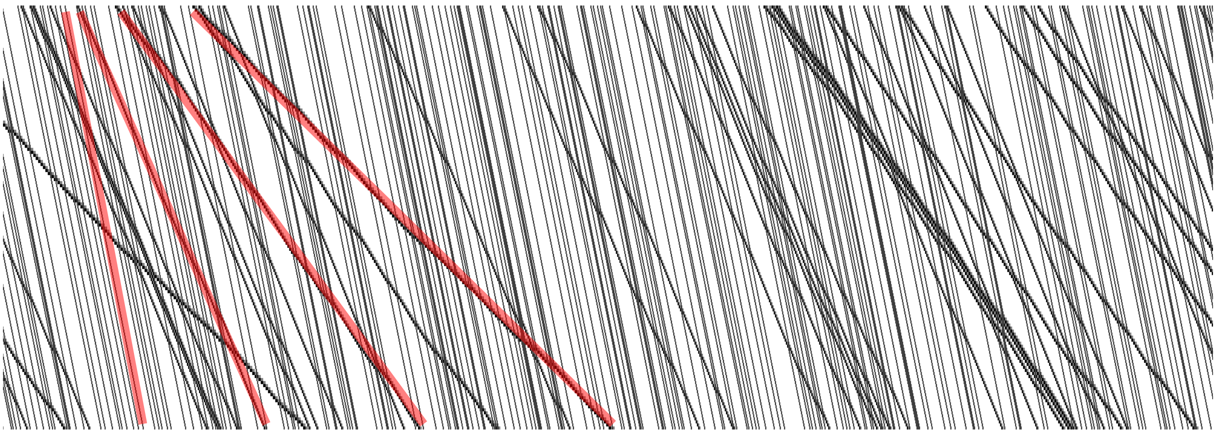

Figure 1.2: Simulation for an i.i.d. configuration with density 0.15. The transparent red lines have deterministic slopes computed by Theorem 1.2, which have been manually shifted so that they would overlay a soliton. This window covers 2000 sites and 140 time steps going downwards, and has been stretched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

given by $M_k\eta(i) :=$ number of the k-solitons appended to the i-th k-slot. Let $\widehat{\mu} = \mu(\,\cdot\,|\,0 \in R\eta)$.

Theorem 1.3. *Let $\zeta = (\zeta_k)_{k\geqslant 1}$ be independent random elements of $(\mathbb{Z}_+)^{\mathbb{Z}}$ with shift-invariant distributions satisfying $\sum_k kE[\zeta_k(0)] < \infty$. Then there exists a unique shift-invariant probability μ on $\mathcal{X}$ such that $M_k\eta \overset{d}{=} \zeta_k$ under $\widehat{\mu}$, and 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. This is in contrast with many stochastic systems, which have at most one invariant measure for each density λ. Moreover, given a sequence $(\rho_k)_k$ specifying the density of k-solitons for each k, there is still a large family of mutually singular states, all having the same specified densities.

The extra assumption needed in Theorem 1.2 is that μ be of the above form, i.e., under $\widehat{\mu}$ each k-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 §4.4.

1.3 Dynamics of components, main tools and paper outline

The two previous results rely heavily on the analysis of how the k-components evolve under T. Then main result in this analysis is Theorem 2.1, which says that ζ_k is translated by k plus a correction caused by the relabeling of slots when larger solitons jump over the origin. For brevity we defer a precise statement to §2.2 where the slot decomposition will have been described in more details.

Another tool that we use abundantly are Palm transformations, which allow us to go back and forth from a configuration seen from an external observer, a typical record or a typical k-soliton. Combining Palm theory with the operator T permits the analysis of the dynamics as seen from a tagged record or a tagged k-soliton. This framework is crucial in the explicit construction of invariant measures and in the analysis of asymptotic speeds.

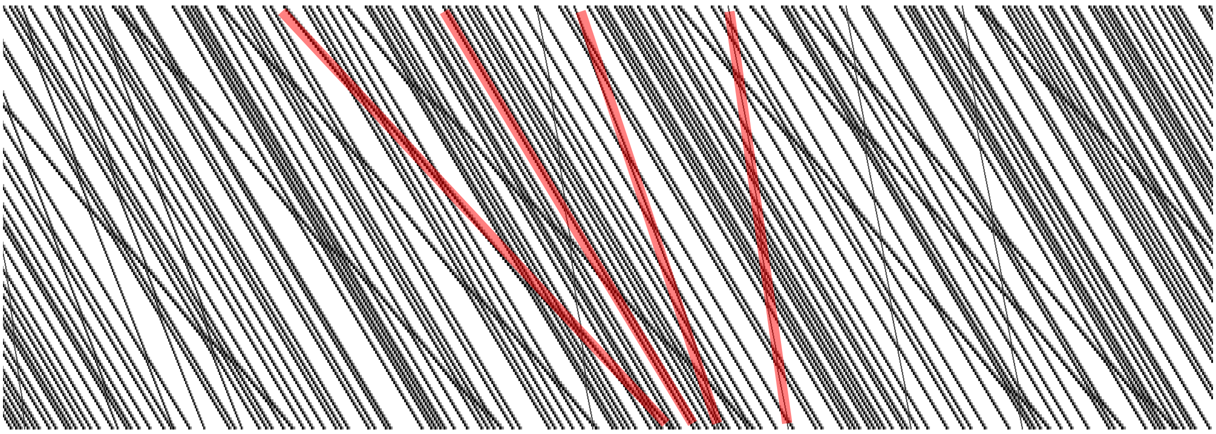

Figure 1.3: Simulation for $(\rho_k)_k = (.006, .005, .1, .003, 0, 0, 0, \dots)$. The initial configuration was obtained by first appending one k-soliton with probability ρ_k after each record, and then applying T a number of times in order to mix. As in Fig. 1.2 it is a 2000x140 window stretched by 5, and red lines are deterministic. (high resolution, color online)

The paper is organized as follows. In §2 we describe the TS algorithm to identify solitons and show that solitons are conserved by the dynamics. We then introduce the slot decomposition of ball configurations, describe how k-components are translated by the dynamics, and finally show that a configuration can be reconstructed from the k-components. In §3 we prove Theorem 1.3 by giving an explicit construction of T-invariant measures which are shift-invariant. In §4 we study the asymptotic speed of tagged solitons to prove Theorems 1.1 and 1.2 and also study the speeds in terms of records. In §5 we complete some proofs postponed in previous sections.

2 Slot decomposition and reconstruction

We define $\xi = \xi[\eta]$ as 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,

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

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 $z < x$. Let $\mathcal{W}_\lambda = \{\xi[\eta] : \eta \in \mathcal{X}_\lambda\}$ and $\mathcal{W} = \{\xi[\eta] : \eta \in \mathcal{X}\}$. Then every $\xi \in \mathcal{W}$

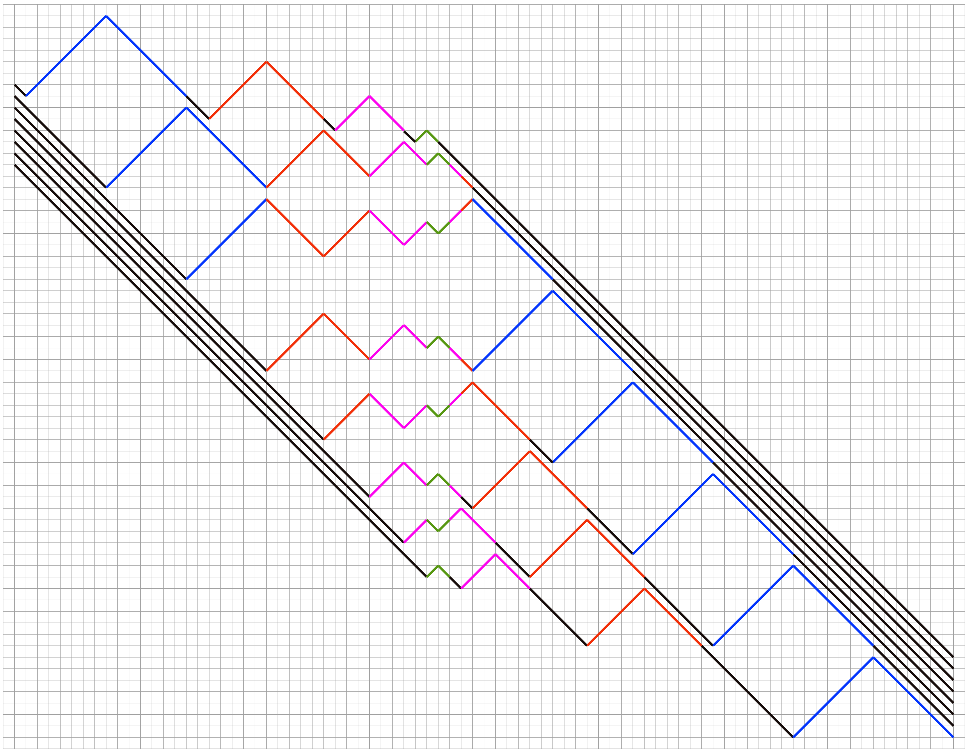

Figure 2.1: Time-evolution of a walk under seven iterations of T. This example has four solitons, of size 7, 5, 3 and 1. Different colors are used to highlighting their conservation. To facilitate view we have shifted the walk at time t by t units down. (color online)

satisfies $\min_{y \leqslant x} \xi(y) \in \mathbb{Z}$ for all $x \in \mathbb{Z}$, and we define

$$T\xi(x) := 2 \min_{y \leqslant x} \xi(y) - \xi(x) = \big[\min_{y \leqslant x} \xi(y)\big] - \Big[\xi(x) - \min_{y \leqslant x} \xi(y)\Big]. \tag{2.1}$$

This amounts to reflecting the walk ξ with respect to its running minimum. The operation T on $\mathcal{W}$ and conservation of solitons are illustrated in Fig. 2.1.

One can see ξ as a *lift* of η which includes an arbitrary choice of vertical shift, or equivalently an arbitrary labeling 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:

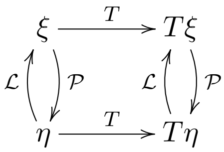

In general this diagram commutes except that the lifting $\mathcal{L}$ misses uniqueness while the projection $\mathcal{P}$ cancels such non-uniqueness. They are analogous to the derivative and indefinite integral where the latter comes with an indeterminate additive constant. If a property is insensitive 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, the above definition of record coincides with the one given at the Introduction. Likewise, (1.2) is equivalent to (2.1), and from the latter one immediately gets Lemma 1.3. Some of the objects considered in this section do however depend on the lift ξ. Note that properties of η always translate to ξ, for instance $\Gamma_m \xi$ means simply $\Gamma_m \eta[\xi]$, etc.

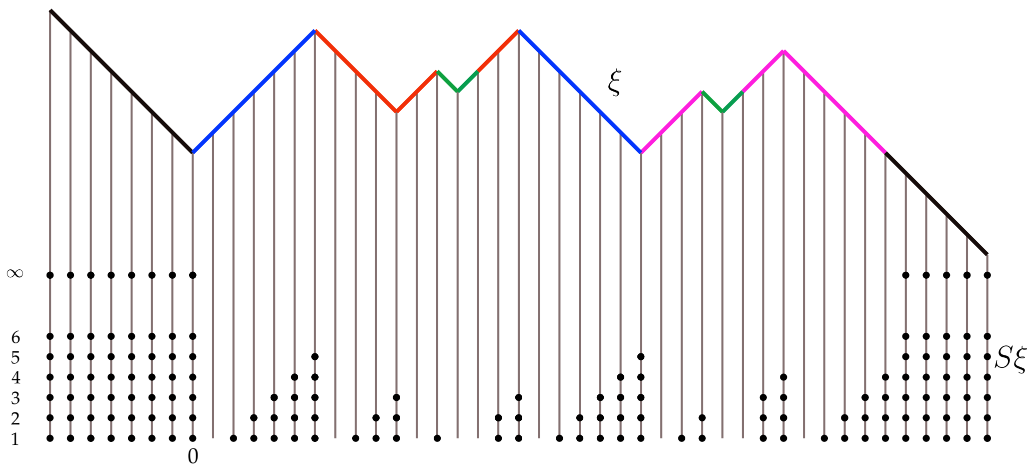

Figure 2.2: Slot configuration of a walk ξ. Different colors correspond to different solitons; records are painted in black. For each site, the number of dots below it indicates its level in slots: there being k dots means a k-slot. (color online)

2.1 Slots and components

We now describe how solitons can be nested inside each other via what we call *slots*. Let γ be a k-soliton. We label the sites in the head and tail of γ in increasing oder by $\mathcal{H}(\gamma) = \{\mathcal{H}_1(\gamma), \dots, \mathcal{H}_k(\gamma)\}$ and by $\mathcal{T}(\gamma) = \{\mathcal{T}_1(\gamma), \dots, \mathcal{T}_k(\gamma)\}$. Given a walk representation ξ, the *slot configuration* $S\xi : \mathbb{Z} \to \{0, 1, 2, \dots\} \cup \{\infty\}$ is defined by

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

For each $k \geqslant 1$ we say that x is a *k-slot* for ξ if $S\xi(x) \geqslant k$. Note that a record is a k-slot for all k, and an m-soliton contains a number $2m - 2k$ of k-slots, see Fig. 2.2. Since 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

$$r(\xi, j) := \min\{x \in \mathbb{Z} : \xi(x) = -j\}. \tag{2.2}$$

This is the leftmost site where the walk ξ takes the value $-j$. If $\xi \in \mathcal{W}$, we have $r(\xi, j) \in \mathbb{Z}$ is well-defined for all $j \in \mathbb{Z}$.

Denote the set of the k-slots of ξ by $\{s_k(\xi, i)\}_{i \in \mathbb{Z}}$, so that the labeling is increasing in i and the 0-th k-slot is at Record 0, that is, $s_k(\xi, 0) = r(\xi, 0)$. We say that a k-soliton γ is *appended* to the i-th k-slot if γ is contained in the interval $[s_k(\xi, i) + 1, s_k(\xi, i+1) - 1]$. Each k-soliton in ξ is appended to a unique k-slot, but any finite number of k-solitons can be appended to a single

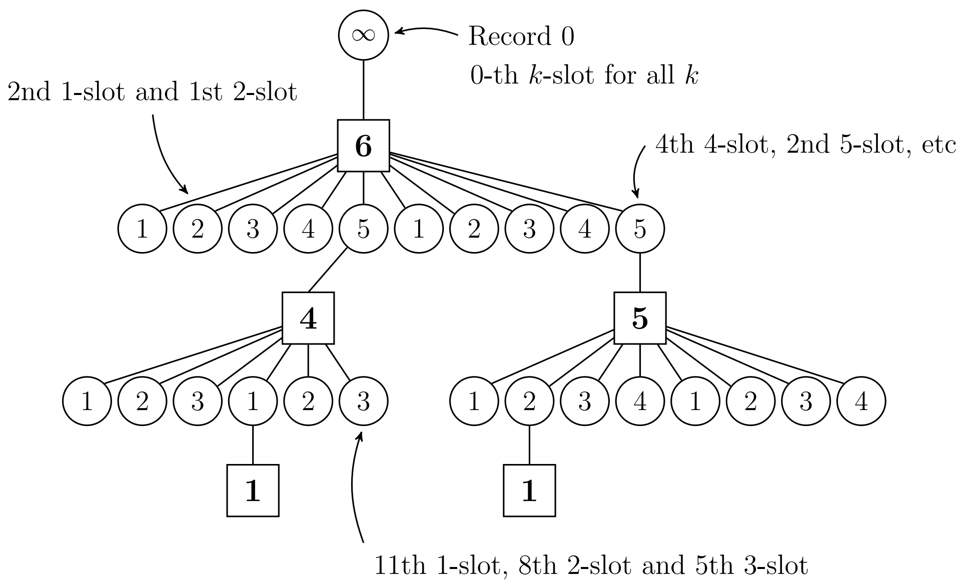

Figure 2.3: An illustration of how the solitons are nested inside bigger solitons via slots, in the same sample configuration as in Fig. 2.2. Solitons are represented by squares and slots by circles. For each $k \geqslant 1$, each slot with index $m \geqslant k$ is a k-slot. We say it is the n-th k-slot, where the ordinal n is determined by counting how many k-slots appear before it in the depth-first order, and the counting starts from the 0-th k-slot present at Record 0.

k-slot. See Fig. 2.3 for how the solitons are nested inside each other via slots.

For $\xi \in \mathcal{X}$ we define the *k-component* of ξ as the configuration $M_k\xi$ of k-solitons appended to the k-slots, that is, $M_k\xi(i)$ denotes the number of the k-solitons appended to the i-th k-slot. In the example of Fig. 2.2, $M_6\xi(0) = 1,\ M_5\xi(2) = 1,\ M_4\xi(2) = 1,\ M_1\xi(9) = 1,\ M_1\xi(18) = 1$, and $M_k\xi(i) = 0$ otherwise.

2.2 Evolution of components

Recall that we can track a tagged soliton γ after t iterations of T by (1.5). Tagged records are also tracked by the formula (2.2). In order to track slots we introduce the flows of solitons and slots, as follows.

Let $\xi \in \mathcal{W}$ and define the *flow of m-solitons through Record 0 by time t* as the number of m-solitons γ to the left of Record 0 in ξ and such that γ^t is to the right of Record 0 in $T^t\xi$:

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

We now define an observable $o^t_k(\xi)$ which 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

$$o^t_k(\xi) := \sum_{m>k} 2(m-k)J^t_m\xi.$$

Using this observable, we define the *relabeled 0-th k-slot at time t* by

$$s^t_k(\xi,0) := s_k(T^t\xi, o^t_k(\xi)),$$

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

$$s^t_k(\xi,i) := s_k(T^t\xi, o^t_k(\xi) + i).$$

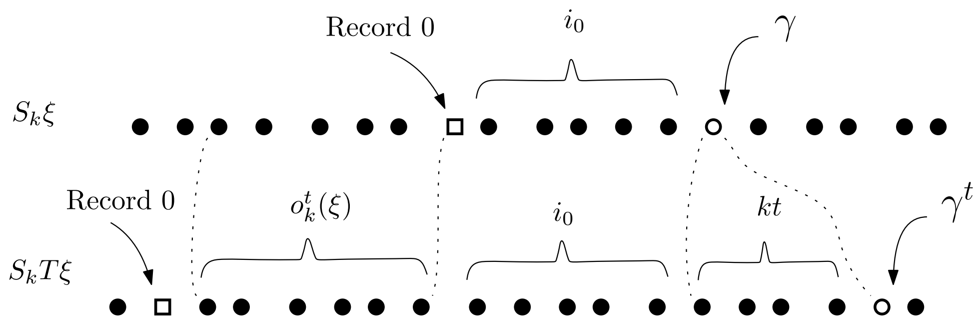

Figure 2.4: Relabeling and tagging k-slots throughout the dynamics. In general the relabeled k-slot will not appear at the same position, and they are identified in M_k by counting from Record 0. But some m-solitons with $m > k$ may cross Record 0, bringing a number $o^t_k(\xi)$ of slots with them, thus the need to relabel the k-slots. The dotted lines mark three relabeled slots and one tagged slot, which happens to contain a k-soliton.

Theorem 2.1. *Let $\xi \in \mathcal{W}$ and $t \in \{1, 2, 3, \dots\}$. Then, for any k-soliton γ of ξ,*

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

that is, the right-to-left flow of relabeled k-slots through a tagged k-soliton between times 0 and t is exactly kt. 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 the label at time t of the relabeled k-slot initially at Record 0. Moreover, for each $k \in \mathbb{N}$, the offset $o_k^t(\xi)$ is determined by $(M_m\xi)_{m>k}$.

Let $\pi \in S_k\xi \subseteq \mathbb{Z}$. Then $\pi = s_k(\xi, j)$ for some j and we define

$$\pi^{k,t} := s_k(T^t\xi, o_k^t(\xi) + kt + j). \tag{2.3}$$

If there is a k-soliton γ appended to the k-slot π of ξ, by Theorem 2.1 the k-soliton γ^t will be appended to the k-slot $\pi^{k,t}$ of $T^t\xi$. In particular, the definition of $\pi^{k,t}$ depends only on $\eta[\xi]$. We note that the difference between the relabeled slots introduced before and tagged slots introduced just now is the factor of kt related to the motion of k-solitons. So a tagged k-soliton crosses k relabeled k-slots per unit time whereas it does not cross tagged k-slots (in fact it just

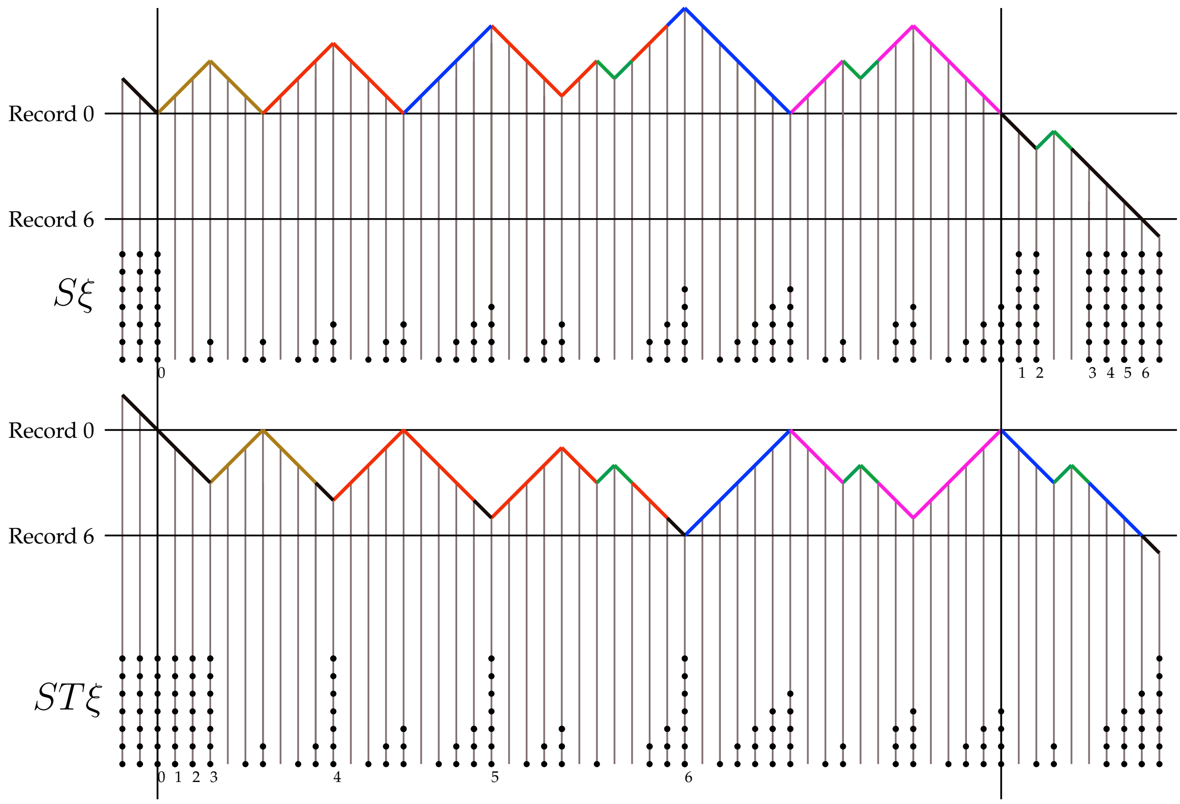

Figure 2.5: We depict $T\xi$ below ξ. The support of excursion ε of ξ is between the two vertical lines. Records 1 to 6 (numbered sites in the picture) are to the right of $\{\varepsilon\}$ in ξ and in positions $a_1(\varepsilon)$ to $a_6(\varepsilon)$ in $T\xi$. The blue 6-soliton is $\tilde{\gamma}$, and $\mathcal{T}(\tilde{\gamma}^1)$, its tail in $T\xi$, is to the right of $\{\varepsilon\}$. (color online)

follows one of them), see Fig. 2.4.

In the remainder of this subsection we prove Theorem 2.1.

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

To show the first statement, it suffices to prove that the number of relabeled k-slots to the right of a k-soliton γ in ξ and to the left of γ^1 in $T\xi$ is exactly k. Consider an excursion ε of height $m \geqslant k$ and let $\tilde{\gamma}$ be the rightmost m-soliton in ε.

Assume first that there are no solitons to the left of ε. Denote

$$a_i(\varepsilon) := \min\{\mathcal{H}_i(\gamma) : \gamma \text{ is a soliton of size} \geqslant i \text{ and contained in } \varepsilon\}, \quad i = 1, \dots, m$$

The slot configuration in the sites $\{\varepsilon\}$ in ξ is modified in $T\xi$ as follows.

$$\begin{aligned} S\xi(a_i(\varepsilon)) &= i - 1, \quad i = 1, \dots, m \\ ST\xi(a_i(\varepsilon)) &= \infty, \quad i = 1, \dots, m. \end{aligned}$$

Indeed, m records that were to the right of $\{\varepsilon\}$ in ξ go respectively to $\{a_1(\varepsilon), \dots, a_m(\varepsilon)\}$ in $T\xi$, and the remaining sites of $\{\varepsilon\}$ keep the same slot configuration:

$$S_k T\xi(x) = S_k\xi(x), \quad \text{for } x \in \{\varepsilon\} \setminus \{a_1(\varepsilon), \dots, a_m(\varepsilon)\},$$

see Fig. 2.5. As a consequence, we have that for each $i \in \{1, \dots, m\}$ and $j \geqslant i$, there is one j-slot "jumping" from the right of $\{\varepsilon\}$ to $a_i(\varepsilon)$. More precisely, one 1-slot goes from the right of $\{\varepsilon\}$ to $a_1(\varepsilon)$, two 2-slots go to $a_1(\varepsilon)$ and $a_2(\varepsilon)$, and so on. Since the 1-solitons of ε are appended to 1-slots in $T\xi$ starting from $a_1(\varepsilon)$, we have that exactly one 1-slot crossed from the right to the left of each 1-soliton in ε. In the same way, each k-soliton of ε is appended to a k-slot of $T\xi$ which is at some $x \geqslant a_k(\varepsilon)$, so the flow of k-slots across any k-soliton in ε is exactly k, concluding the proof.

We now drop the assumption that there are no solitons to the left of ε. Suppose there is an excursion ε' of height $n \geqslant 2$ to the left of the excursion ε and there are $\ell < n$ records between ε' and ε (if there are more than n records, then ε' will not interfere with ε). Then, the records that in the previous case were going to $a_i(\varepsilon)$ are now going to $a_{n-\ell-i}(\varepsilon')$ in $T\xi$ for $i \leqslant n - \ell$, while the slot configuration of ξ in $a_{n-\ell-i}(\varepsilon')$ goes to $a_i(\varepsilon)$ in $T\xi$. Hence, as before the flow of relabeled k-slots across any k-soliton in ε equals k. If there are more than two such excursions, the description gets more complicated but the principle remains the same.

For the third statement, we need to show that $o_k^t(\xi) = o_k^t(\xi')$ whenever $\xi, \xi' \in \widehat{\mathcal{W}}$ are such that $(M_m\xi : m > k) = (M_m\xi' : m > k)$, that is, we need to show that for any $m > k$ the left-to-right flow of m-solitons through Record 0 for ξ and ξ' coincide, which means $J_m^1\xi = J_m^1\xi'$. To prove that, we observe that $J_m^1\xi$ is given as follows. Let $k_0 = k_0(\xi) = \max\{k : M_k\xi(j) \neq$

0 for some $j = -1, \dots, -k\}$. Then $J_k^1(\xi) = 0$ and thus $o_k^1(\xi) = 0$ for all $k \geqslant k_0$. For $k = k_0 - 1$, since $J_m(\xi) = 0$ for all $m > k$, we have $o_k^1(\xi) = 0$ and $J_k^1(\xi)$ is determined by $M_k\xi$ only. For $k = k_0 - 2$, we have $o_k^1(\xi)$ determined by $J_m^1(\xi)$ for $m > k$, which in turn is determined by $(M_m(\xi) : m > k)$ and hence $J_k^1(\xi)$ is determined by $(M_m\xi : m > k)$. 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 2.1.

2.3 Reconstructing the configuration from the components

Consider the map $\xi \mapsto M\xi := (M_k\xi)_{k\geqslant 1}$ and let

$$\mathcal{M} := \{M\xi : \xi \in \mathcal{W}\} \subseteq \left((\mathbb{N}_0)^{\mathbb{Z}}\right)^{\mathbb{N}}$$

denote the set of possible component sequences obtained from decomposing configurations in $\mathcal{W}$. Since the decomposition $\xi \mapsto M\xi$ is insensitive to horizontal shifts, it is not possible to determine ξ knowing $(M_k\xi)_{k\geqslant 1}$. So we define the space

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

We remark that, unlike the lift $\xi[\eta]$ from $\mathcal{X}$ to $\mathcal{W}$ which was not unique, for η in

$$\widehat{\mathcal{X}} := \{\eta \in \mathcal{X} : 0 \text{ is a record for } \eta\}$$

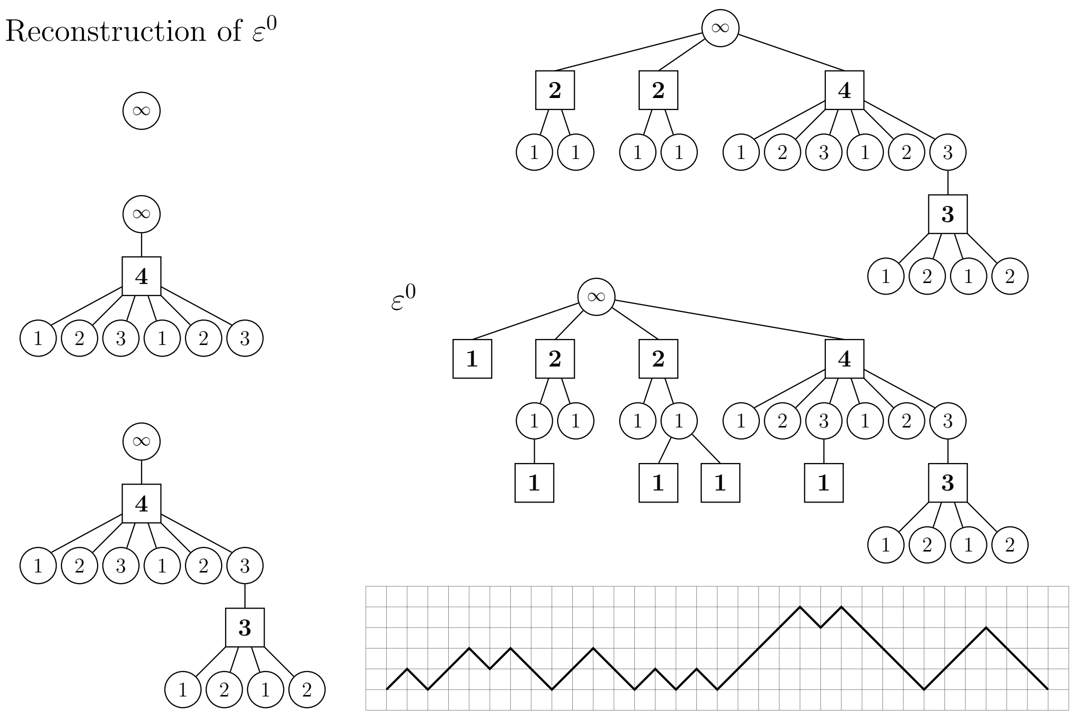

Figure 2.6: Reconstruction algorithm for a single excursion. This example is obtained using the field ζ shown in Fig. 2.7.

there is a unique lift $\xi[\eta]$ which is in $\widehat{\mathcal{W}}$, we denote it $\xi^\circ[\eta]$. So the operator M defined in the Introduction is given by $M_k\eta = M_k\xi^\circ[\eta]$. To remain consistent with the previous subsections, we continue working with ξ instead of η.

In the sequel we show that the map M restricted to $\widehat{\mathcal{W}}$ is invertible.

The *height* of the excursion between Records 0 and 1 is defined as 0 if the excursion is empty or as $\max\{k \geqslant 1 : M_k\xi(0) > 0\}$ otherwise. In general, denoting $i_k(j)$ the label of the k-slot located at Record j, we define the height of the excursion between Record j and Record $j+1$ by

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

which is well-defined for every $\xi \in \mathcal{W}$. Both $m(j)$ and $i_k(j)$ depend on ξ but we omit it in the notation. Let

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

Note that, for each $\xi \in \mathcal{W}$ and $j \in \mathbb{Z}$, $i_k(j) = r(\xi, j)$ for all large k. Since $m(j)$ is finite, we have that

$$\text{for each } i \in \mathbb{Z},\ \zeta_k(i) = 0 \text{ for all large } k. \tag{2.4}$$

Finally, define the *support* of an excursion ε between successive records $y_1 < y_2$ by $\{\varepsilon\} := \{y_1 + 1, \dots, y_2 - 1\}$ and denote the number of k slots in the excursion by

$$n_k(\varepsilon) := 1 + |S_k\xi \cap \{\varepsilon\}|, \tag{2.5}$$

where the term 1 refers to the record y_1 preceding $\{\varepsilon\}$ and the second term counts the number of k-slots belonging to an m-soliton of ε for some $m > k$.

Let $\zeta = (\zeta_k)_{k\geqslant 1} \in \mathcal{M}$, so in particular ζ satisfies (2.4). We first give an algorithm which permits

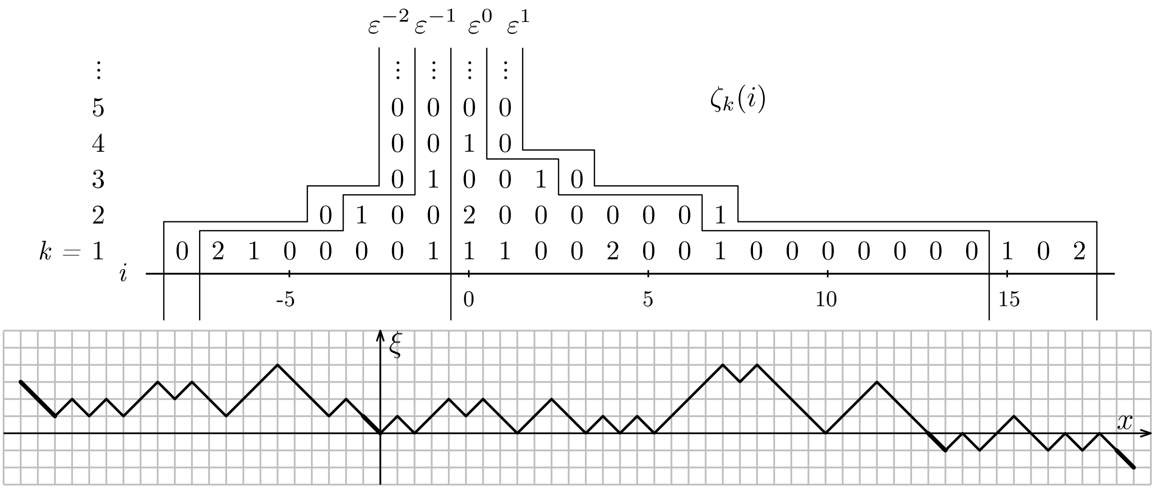

Figure 2.7: Reconstruction of ξ from ζ. In the lower part we show Records -2 to 2 in boldface and the excursions between them. Above we show the parts of the field ζ that used in the reconstruction of $\varepsilon^{-2}, \varepsilon^{-1}, \varepsilon^0, \varepsilon^1$. Reconstruction of ε^0 was shown in Fig. 2.6 and $\varepsilon^1, \varepsilon^{-1}, \varepsilon^{-2}$ is shown in Fig. 2.8.

to reconstruct the excursion ε of ξ between Records 0 and 1. Here is the algorithm:

Let $\varepsilon \in \widehat{\mathcal{W}}$ denote the empty excursion
Let $m := \min\{k \geqslant 0 : \zeta_{k'}(0) = 0 \text{ for all } k' > k\}$
for $k = m, m-1, \dots, 2, 1$ **do**
  Let $n_k := n_k(\varepsilon) = \#\{x \in S_k\varepsilon : r(\varepsilon, 0) \leqslant x < r(\varepsilon, 1)\}$, as in (2.5)
  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

Note that m is well-defined by (2.4). In case $m = 0$, the algorithm produces an empty excursion. The procedure is illustrated in Fig. 2.6. Call ε^0 the excursion just constructed. Construct ε^1, the excursion between Records 1 and 2, using the same algorithm but with the data $\zeta^1 = (\zeta^1_k)_{k\geqslant 1}$, where each component is given by

$$\zeta^1_k = \big(\zeta_k(n_k + i)\big)_{i\geqslant 0},$$

which consists of the entries of ζ with non-negative indices i not used in the reconstruction of ε^0. Note that ζ^1 also satisfies (2.4). Iterate this procedure to construct an infinite sequence of excursions $(\varepsilon^j)_{j=0,1,2,\dots}$. See Fig. 2.8.

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 $\zeta^{-1} = (\zeta^{-1}_k)_{k\geqslant 1}$ where each component is given by $\zeta^{-1}_k = \big(\zeta_k(i)\big)_{i<0}$ and use ζ^{-1} to construct ε^{-1}. Then define $\zeta^{-2} = (\zeta^{-2}_k)_{k\geqslant 1}$ where each component is given by $\zeta^{-2}_k = \big(\zeta_k(i - n_k)\big)_{i<0}$ and use it to construct ε^{-2}. Iterate this procedure to construct an infinite sequence of excursions $(\varepsilon^j)_{j=-1,-2,\dots}$.

Put Record 0 at the origin and concatenate the excursions with one record between each pair

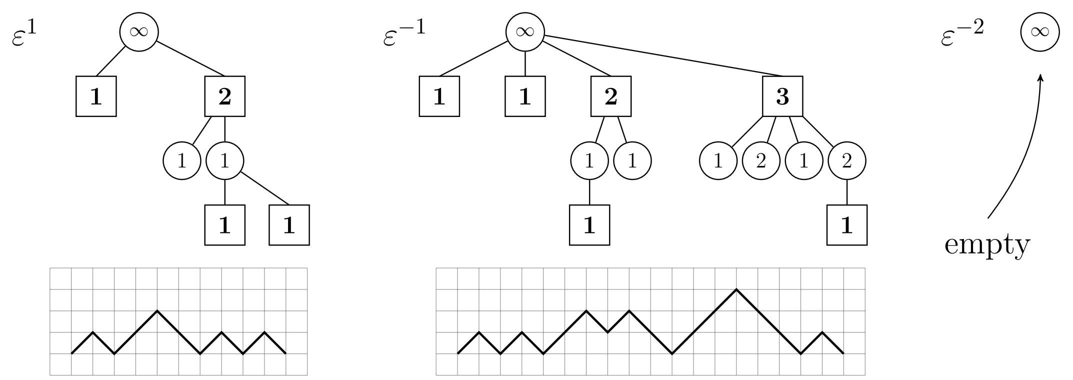

Figure 2.8: Reconstruction algorithm for other excursions. The procedure is the same as in Fig. 2.6 but all the intermediate steps are omitted.

of consecutive excursions. This yields a walk denoted ξ^*, shown in Fig. 2.7.

Call $M^{-1} : \zeta \mapsto \xi^*$ the resulting transformation. We claim that M^{-1} is the inverse map of M restricted to $\widehat{\mathcal{W}}$, that is, $M^{-1}M\xi = \xi$ for $\xi \in \widehat{\mathcal{W}}$. We will sketch the proof of a single excursion ε omitting tedious details. For $k \geqslant 0$, denote by $\varepsilon_{[k]}$ the ball configuration obtained by removing all the boxes belonging to an ℓ-soliton with $\ell \leqslant k$. 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 $m > k$, as the 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 $\varepsilon_{[k+1]}$, hence it correctly finds ε, completing the proof.

3 Invariant measures

In this section we prove Theorem 1.3. More precisely, we show explicitly how invariant measures can be constructed by specifying the distribution of each k-component ζ_k.

We refer to probability measures as simply *measures*. We also refer to measurable functions as *random elements*, and refer to the push-forward of a pre-specified measure by such functions as the *law* of these random elements. A measure μ on $\mathcal{X}$ is *T-invariant* if $\mu \circ T^{-1} = \mu$. The following is proved in §5.

Proposition 3.1. *For $\lambda < \frac{1}{2}$, the product measure* $\text{Bernoulli}(\lambda)^{\otimes\mathbb{Z}}$ *is T-invariant.*

Besides the i.i.d. states, 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.

3.1 Construction of the measures

Our recipe to produce invariant measures uses the construction described in §2.3, which gives a distribution $\widehat{\mu}$ of configurations "seen from a typical record." The proof of Theorem 1.3 is based on properties of $\widehat{\mu}$ and how they relate to the the dynamics. This relationship is given by the Palm theory, which we briefly recall now.

Denote the shifts by

$$\theta\eta(y) := \eta(y+1) \text{ for } \eta \in \{0,1\}^{\mathbb{Z}}, \quad \theta A = A - 1 \text{ for } A \subseteq \mathbb{Z}.$$

Let $\eta \in \{0,1\}^{\mathbb{Z}}$. Recalling (1.1), note that $\theta R\eta = R\theta\eta$. Suppose $\eta \in \widehat{\mathcal{X}}$, and denote

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

If such η is random with law $\widehat{\mu}$ such that $\widehat{\mu}(\mathtt{r}) := \int \mathtt{r}(\eta)\, \widehat{\mu}(\mathrm{d}\eta) < \infty$, its *inverse-Palm measure*

$\mu = \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu})$ is defined as follows. For every test function φ,

$$\int \varphi(\eta)\,\mu(\mathrm{d}\eta) := \int \frac{\sum_{i=1}^{\mathtt{r}(\eta)} \varphi(\theta^i \eta)}{\mathtt{r}(\eta)} \cdot \frac{\mathtt{r}(\eta)}{\widehat{\mu}(\mathtt{r})} \widehat{\mu}(\mathrm{d}\eta). \tag{3.2}$$

In words, 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, and then choose a site uniformly from this excursion to place the origin.

We define a bijection between $R\eta$ and $RT\eta$ as follows. For $x \in R\eta$, let $\psi_\eta(x) := r(T\xi, j)$, where $x = r(\xi, j)$. Note that this definition does not depend on the lift $\xi[\eta]$. We define the *record-shift* $\widehat{\theta} : \widehat{\mathcal{X}} \to \widehat{\mathcal{X}}$ and the *dynamics seen from a record* $\widehat{T} : \widehat{\mathcal{X}} \to \widehat{\mathcal{X}}$ by

$$\widehat{\theta}\eta := \theta^{\mathtt{r}(\eta)}\eta \quad \text{and} \quad \widehat{T}\eta := \theta^{\psi_\eta(0)} T\eta.$$

The following are standard properties of Palm measures which we prove in §3.4.

Lemma 3.3. *If $\widehat{\mu}$ is $\widehat{\theta}$-invariant and $\widehat{\mu}(\mathtt{r}) < \infty$, then* $\mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu})$ *is θ-invariant, supported on $\mathcal{X}$ and satisfies* $\mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu}) \circ T^{-1} = \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu} \circ \widehat{T}^{-1})$.

Lemma 3.4. *If moreover $\widehat{\mu}$ is $\widehat{\theta}$-ergodic, then* $\mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu})$ *is θ-ergodic.*

Let $\zeta = (\zeta_k)_{k \geqslant 1}$ be independent random elements of $(\mathbb{Z}_+)^{\mathbb{Z}}$ with shift-invariant distribution whose expectations satisfy $\sum_k kE[\zeta_k(0)] < \infty$. Take $\xi = M^{-1}\zeta$ as the configuration reconstructed from soliton components ζ according to the algorithm described in §2.3 and depicted in Fig. 2.7. Let $\widehat{\mu}$ denote the resulting law of $\eta = \eta[\xi]$.

Proposition 3.5. *The measure $\widehat{\mu}$ defined above is $\widehat{T}$-invariant and $\widehat{\theta}$-invariant, and it also satisfies $\widehat{\mu}(\mathtt{r}) < \infty$. 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 1.3.

Let $\widehat{\mu}$ be the law of $\eta = M^{-1}\zeta \in \widehat{\mathcal{X}}$, and define

$$\mu := \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu}).$$

By Proposition 3.5 and Lemma 3.3, μ is θ-invariant and supported on $\mathcal{X}$. Also,

$$\mu \circ T^{-1} = \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu} \circ \widehat{T}^{-1}) = \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu}) = \mu,$$

so μ is also T-invariant. Moreover, under the i.i.d. assumption, the second part of Proposition 3.5 combined with Lemma 3.4 imply that μ is θ-ergodic.

Remark 3.6 (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 [FG18].

Remark 3.7 (Counter-examples)**.** It is possible for law μ of η to be θ-ergodic and T-invariant while its components $M_k\eta$ not being 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 $\zeta = (\zeta_k)_{k\geqslant 1}$. The reader can check that this example satisfies the stated properties. Likewise, 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 $\zeta = (\zeta_k)_{k\geqslant 1}$. We conjecture that if μ is T-invariant and θ-mixing then the ζ_k are independent over k and each one is θ-mixing.

3.2 Invariance of the reconstructed configuration

We now prove the main part of Proposition 3.5, 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 $\widehat{\mu}(\mathtt{r}) < \infty$ is given in §3.3. Denote by E the integral with respect to the law of ζ, and by $\mathcal{F}(\cdot)$ the sigma-field generated by the random elements $(\cdot)$. First note that $\xi := M^{-1}\zeta$ is a.s. well-defined. Indeed, since $\sum_k E[\zeta_k(0)] < \infty$, the random field ζ a.s. satisfies (2.4) by Borel-Cantelli.

To show that $\widehat{\mu}$ is $\widehat{T}$-invariant it suffices to show that the slot decomposition of $\widehat{T}\eta$ has the same law as ζ. More precisely, it suffices to show

$$E\Big(\prod_{k=1}^{n} \varphi_k(M_k\widehat{T}\eta)\Big) = \prod_{k=1}^{n} E\varphi_k(\zeta_k), \qquad \text{for each } n \geqslant 1, \tag{3.8}$$

for test functions $\varphi_1, \dots, \varphi_n$, $k \geqslant 1$.

Note that we can write

$$\begin{aligned}
&E\big(\varphi_k(M_k\widehat{T}\eta) \,\big|\, \mathcal{F}(\zeta_m : m > k)\big) \\
&\quad = E\big(\varphi_k(\theta^{-o_k^1(\xi^\circ[\eta])-k} M_k\eta) \,\big|\, \mathcal{F}(\zeta_m : m > k)\big) \qquad \text{(by Theorem 2.1)} \\
&\quad = E\big(\varphi_k(\theta^{-o_k^1(\xi^\circ[\eta])-k} \zeta_k) \,\big|\, \mathcal{F}(\zeta_m : m > k)\big) \qquad \text{(because } M_k\eta = \zeta_k) \\
&\quad = E\varphi_k(\zeta_k),
\end{aligned}$$

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 (3.8) is then

$$\begin{aligned}
E\Big(\prod_{i=k}^{n} \varphi_i(M_i\widehat{T}\eta)\Big) &= E\Big(E\Big(\prod_{i=k}^{n} \varphi_i(M_i\widehat{T}\eta)\Big|\mathcal{F}(\zeta_m : m > k)\Big)\Big), \\
&= E\Big(\prod_{i=k+1}^{n} \varphi_i(M_i\widehat{T}\eta) E\Big(\varphi_k(M_k\widehat{T}\eta)\Big|\mathcal{F}(\zeta_m : m > k)\Big)\Big), \\
&= E\varphi_k(\zeta_k)\; E\Big(\prod_{i=k+1}^{n} \varphi_i(M_i\widehat{T}\eta)\Big);
\end{aligned}$$

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

Finally, consider the transformation $M^{-1} : \zeta \mapsto \eta^*$, defined in §2.3 and call ε_*^0 the excursion of

η^* between Records 0 and 1. The construction of §2.3 gives

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

So it suffices to show that $(\theta^{n_k(\varepsilon_*^0)}\zeta_k)_{k\geqslant 1}$ has the same law as $(\zeta_k)_{k\geqslant 1}$. But $n_k(\varepsilon_*^0)$ is determined by $(\zeta_m)_{m>k}$, thus 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 $\widehat{\mu}$ is $\widehat{\theta}$-invariant.

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.

3.3 Expected excursion length

We continue the proof of Proposition 3.5 proving that $\widehat{\mu}(\mathtt{r}) < \infty$. The proof is probabilistic but it could be reformulated in terms of the spectrum of an infinite sub-Markovian matrix. We start by showing that the system

$$w_k = 1 + \sum_{m>k} 2(m-k)w_m\alpha_m, \qquad k = 0, 1, 2, \dots \tag{3.9}$$

has a unique finite solution $w = (w_k)_{k\geqslant 0}$, where

$$\alpha_k := E[\zeta_k(0)]. \tag{3.10}$$

Then we will 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

$$\rho_k = \alpha_k w_k. \tag{3.11}$$

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

$$\widehat{\mu}(\mathtt{r}) = w_0 = 1 + \sum_{m\geqslant 1} 2m\,\rho_m < \infty. \tag{3.12}$$

So we start by studying (3.9). Let $c_k := 2\sum_{m>k}(m-k)\alpha_m$ and take $\tilde{k}$ such that $\sum_{m>\tilde{k}} 4m\alpha_m < 1$, so $c_k < \frac{1}{2}$ for $k \geqslant \tilde{k}$. Let $K := \{k \in \mathbb{N} : k \geqslant \tilde{k}\} \cup \{\aleph\}$ and consider a Markov chain $(X_n)_{n\geqslant 0}$ on K with absorbing state $\aleph$ and transition probabilities $q(k,m) := 2(m-k)\alpha_m \mathbb{1}\{m > k\}$; $q(k,\aleph) = 1 - c_k$; $q(\aleph,\aleph) = 1$ and $q(k,m) = 0$ otherwise. Define the absorption time by $\tau := \inf\{n \geqslant 0 : X_n = \aleph\}$. Denote by E_k the law of $(X_n)_n$ starting from k. By conditioning on X_1, we see that the expectations $w_k = E_k\tau$ satisfy the system (3.9). Since $c_k \geqslant c_{k+1}$, we have $P_k(\tau > n) \leqslant c_k^n$ and thus $w_k = E_k\tau \leqslant \frac{1}{1-c_k} < 2 < \infty$, for $k \geqslant \tilde{k}$. Since $w_{\tilde{k}} < \infty$, using (3.9) with $k = \tilde{k} - 1$ we get $w_{\tilde{k}-1} < \infty$, and iterating this argument we get $w_k < \infty$ for all k.

We now consider truncated approximations for the reconstruction algorithm of §2.3. Let

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

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 (2.5). 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 $n \to \infty$. On the other hand, since 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 $\alpha_m^n := \alpha_m \mathbb{1}_{m\leqslant n}$ denote the expected number of m-solitons per m-slot in $\zeta^{[n]}$. Since $W_k^n\zeta$ is a function of $(\zeta_m : m > k)$ which is independent of ζ_k, the expected number of m-solitons in $\varepsilon^{[n]}$ is $w_m^n \times \alpha_m^n$. Therefore, $(w_k^n)_{k\geqslant 0}$ and $(\alpha_k^n)_{k\geqslant 1}$ satisfy the system (3.9). Finally, since $w_k^n < 2$ for all $k \geqslant \tilde{k}$ and $n \in \mathbb{N}$, w_k is finite for every k and therefore (3.12) is satisfied, concluding the proof.

3.4 Palm transformations

There is a bijection between θ-invariant measures μ on $\mathcal{X}$ and $\widehat{\theta}$-invariant measures $\widehat{\mu}$ on $\widehat{\mathcal{X}}$ with $\widehat{\mu}(\mathtt{r}) < \infty$. Given such a μ, we define $\widehat{\mu} = \mathrm{Palm}_{\mathbb{Z}}^{R}(\mu)$ by

$$\widehat{\mu} := \mu(\,\cdot\,|\, 0 \in R\eta), \tag{3.13}$$

that is, $\widehat{\mu}$ equals μ conditioned on η having a record at $x = 0$.

The two maps $\mathrm{Palm}_{\mathbb{Z}}^{R}$ and $\mathrm{Palm}_{R}^{\mathbb{Z}}$ are the inverse of each other. Although their descriptions (3.2) and (3.13) may look different, they are in fact very similar. Indeed, one can think of the conditioning on $0 \in R\eta$ as biasing by the number of records at the origin. So rewriting the two definitions we get

$$\int \varphi(\eta)\,\widehat{\mu}(\mathrm{d}\eta) = \frac{\int \mathbb{1}_{\widehat{\mathcal{X}}}(\eta)\,\varphi(\eta)\,\mu(\mathrm{d}\eta)}{\int \mathbb{1}_{\widehat{\mathcal{X}}}(\eta)\,\mu(\mathrm{d}\eta)} \text{ and } \int \varphi(\eta)\,\mu(\mathrm{d}\eta) = \frac{\int \sum_{i=1}^{\mathtt{r}(\eta)} \varphi(\theta^i\eta)\,\widehat{\mu}(\mathrm{d}\eta)}{\int \mathtt{r}(\eta)\,\widehat{\mu}(\mathrm{d}\eta)}.$$

These expressions are the discrete version of the *inversion formula* (8.4.14°) on p.264 of [Tho00]. Theorem 8.4.1 on p.260 says that this is a bijection between θ-invariant measures μ on $\mathcal{X}$ and $\widehat{\theta}$-invariant measures $\widehat{\mu}$ on $\widehat{\mathcal{X}}$ with $\widehat{\mu}(\mathtt{r}) < \infty$.

For use in §4, we note that one can conveniently pass from one random subset to another. Let $Z\eta$ be a subset of $\mathbb{Z}$ which depends on η in a translation-covariant way ($Z\theta\eta = \theta Z\eta$) and which has μ-a.s. positive density. Writing $\tilde{\mu} := \mu(\,\cdot \mid 0 \in Z\eta)$, we have

$$\tilde{\mu} = \mathrm{Palm}_{\mathbb{Z}}^{Z}(\mu) = \mathrm{Palm}_{\mathbb{Z}}^{Z}(\mathrm{Palm}_{R}^{\mathbb{Z}}(\widehat{\mu})) = (\mathrm{Palm}_{\mathbb{Z}}^{Z} \circ \mathrm{Palm}_{R}^{\mathbb{Z}})(\widehat{\mu}) =: \mathrm{Palm}_{R}^{Z}(\widehat{\mu}),$$

where the operator $\mathrm{Palm}_R^Z : \widehat{\mu} \mapsto \tilde{\mu}$ is given as follows. Writing $\tilde{\mathcal{X}} := \{\eta \in \mathcal{X} : 0 \in Z\eta\}$, we have

$$\int \varphi(\eta)\,\tilde{\mu}(\mathrm{d}\eta) = \frac{\int \sum_{i=1}^{\mathtt{r}(\eta)} \mathbb{1}_{\tilde{\mathcal{X}}}(\theta^i\eta)\varphi(\theta^i\eta)\,\widehat{\mu}(\mathrm{d}\eta)}{\int \sum_{i=1}^{\mathtt{r}(\eta)} \mathbb{1}_{\tilde{\mathcal{X}}}(\theta^i\eta)\,\widehat{\mu}(\mathrm{d}\eta)}, \tag{3.14}$$

which follows immediately by combining the two equations above, the first with $\tilde{\mu}$ instead of $\widehat{\mu}$ and the second with $\varphi \cdot \mathbb{1}_{\tilde{\mathcal{X}}}$ instead of φ.

Proof of Lemma 3.3. The property saying that μ is θ-invariant and supported on $\mathcal{X}$ has been quoted above, so it remains to show the identity. We include a proof of this classical result [Har71, PS73] in our context for convenience of the reader. Denote $\mu = \mathrm{Palm}_R^{\mathbb{Z}}(\widehat{\mu})$. Since $\mu \mapsto \mathrm{Palm}_{\mathbb{Z}}^R(\mu)$ is a bijection, the identity is equivalent to $\mathrm{Palm}_{\mathbb{Z}}^R(\mu \circ T^{-1}) = \widehat{\mu} \circ \widehat{T}^{-1}$, which we prove now. For a test function φ,

$$\begin{aligned}
\mathrm{Palm}_{\mathbb{Z}}^R(\mu \circ T^{-1})\varphi &= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \int \mathbb{1}_{\widehat{\mathcal{X}}}(T\eta)\varphi(T\eta)\,\mu(\mathrm{d}\eta) \\
&= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \int \sum_x \mathbb{1}_{[x \in R\eta, \psi_\eta(x)=0]}\,\varphi(T\eta)\,\mu(\mathrm{d}\eta) \\
&= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \sum_x \int \mathbb{1}_{[0 \in R\theta^x\eta, \psi_{\theta^x\eta}(0)=-x]}\,\varphi(T\eta)\,\mu(\mathrm{d}\eta) \\
&= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \sum_x \int \mathbb{1}_{[0 \in R\eta, \psi_\eta(0)=-x]}\,\varphi(T\theta^{-x}\eta)\,\mu(\mathrm{d}\eta) \\
&= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \int \mathbb{1}_{[0 \in R\eta]} \sum_x \mathbb{1}_{[\psi_\eta(0)=x]}\,\varphi(T\theta^{x}\eta)\,\mu(\mathrm{d}\eta) \\
&= \tfrac{1}{\mu(\widehat{\mathcal{X}})} \int \mathbb{1}_{[0 \in R\eta]}\varphi(\widehat{T}\eta)\,\mu(\mathrm{d}\eta) \\
&= \int \varphi(\widehat{T}\eta)\,\widehat{\mu}(\mathrm{d}\eta) = (\widehat{\mu} \circ \widehat{T}^{-1})\varphi.
\end{aligned}$$

The forth identity holds by translation invariance of μ. ☐

Proof of Lemma 3.4. It suffices to show that the Cesàro limits are μ-a.s. constant. Let $k_n \in \mathbb{Z}$ be such that $r(\eta, k_n) \leqslant n < r(\eta, k_n + 1)$. Then,

$$\begin{aligned}
\lim_n \frac{1}{n} \sum_{x=1}^{n} \varphi(\theta^x\eta) &= \lim_n \frac{k_n}{n} \frac{1}{k_n} \sum_{i=0}^{k_n-1} \sum_{j=1}^{r(\widehat{\theta}^i\eta,1)} \varphi(\theta^j\widehat{\theta}^i\eta) + \lim_n \frac{1}{n} \sum_{x=r(\eta,k_n)+1}^{n} \varphi(\theta^x\eta) \\
&= \lim_n \frac{\frac{1}{k_n} \sum_{i=0}^{k_n-1} \sum_{j=1}^{r(\widehat{\theta}^i\eta,1)} \varphi(\theta^j\widehat{\theta}^i\eta)}{\frac{1}{k_n} \sum_{i=0}^{k_n-1} r(\widehat{\theta}^i\eta, 1)} \\
&= \frac{\int \sum_{j=1}^{\mathtt{r}(\eta)} \varphi(\theta^j\eta)\,\widehat{\mu}(\mathrm{d}\eta)}{\int \mathtt{r}(\eta)\,\widehat{\mu}(\mathrm{d}\eta)} = \mu\varphi.
\end{aligned}$$

☐

4 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 speed v_k exists and equals the expected length of the jump of a typical k-soliton in one step.

In §4.1 we show that the speed v_k that appears in (1.6) is μ-a.s. well-defined and is given by (4.1). We also show that it is given by (4.2), and in particular it is finite. Analyzing the interaction between solitons of different sizes, in §4.2 we show that the speeds $(v_k)_k$ satisfy (1.7). Finally, in §4.3 analyze the formula (4.2) using the description of μ from §3.1 to show that the speeds are given by (1.8).

4.1 Existence of speeds via Palm measure and ergodicity

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.

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 decomposed as

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

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 $0 \in \Gamma_k^\circ\eta$. Let $\widehat{\mu}_k := \mathrm{Palm}_{\mathbb{Z}}^{\Gamma_k^\circ}(\mu)$ be the Palm measure of μ with respect to $\Gamma_k^\circ\eta \subseteq \mathbb{Z}$, i.e., $\widehat{\mu}_k = \mu(\,\cdot\,|\, 0 \in \Gamma_k^\circ\eta)$. For $\eta \in \widehat{\mathcal{X}}^k$, let $\mathtt{r}^k(\eta) := \inf\{x \geqslant 1 : x \in \Gamma_k^\circ\eta\}$ and define $\widehat{\theta}_k : \widehat{\mathcal{X}}^k \to \widehat{\mathcal{X}}^k$ as the "shift to the next k-soliton" given by $\widehat{\theta}_k\eta := \theta^{\mathtt{r}^k(\eta)}\eta$. 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.

As in Lemma 3.3, T-invariance of μ implies $\widehat{T}_k$-invariance of $\widehat{\mu}_k$. On the other hand, for $\eta \in \widehat{\mathcal{X}}^k$, and γ containing $x = 0$ the above decomposition becomes

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

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

$$v_k := \int \Delta_\eta^k(0)\, \widehat{\mu}_k(\mathrm{d}\eta). \tag{4.1}$$

It remains to show that it is in fact non-random. Consider the field

$$\tilde{v}_k(\eta, z) := \begin{cases} \lim\limits_{t\to\infty} \dfrac{x(\gamma^t)}{t}, & \text{if } z = x(\gamma) \text{ for } \gamma \in \Gamma_k(\eta), \\ 0, & \text{otherwise.} \end{cases}$$

Since $(\tilde{v}_k(\eta, x))_{x\in\mathbb{Z}}$ is a $\widehat{\theta}_k$-covariant function of η, its distribution is $\widehat{\theta}_k$-ergodic. On the other hand, since different tagged k-solitons never overtake each other, we have

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

By $\widehat{\theta}_k$-ergodicity, this implies that the non-zero entries of this field are a.s. constant.

Finally, applying (3.14) to (4.1) we get

$$v_k = \frac{\int \sum_{i=1}^{\mathtt{r}(\eta)} \mathbb{1}_{\widehat{\mathcal{X}}^k}(\theta^i\eta)\Delta_{\theta^i\eta}^k(0)\,\widehat{\mu}(\mathrm{d}\eta)}{\int \sum_{i=1}^{\mathtt{r}(\eta)} \mathbb{1}_{\widehat{\mathcal{X}}^k}(\theta^i\eta)\,\widehat{\mu}(\mathrm{d}\eta)} = \frac{1}{\rho_k}\int \sum_{y=1}^{\mathtt{r}(\eta)} \Delta_\eta^k(y)\,\widehat{\mu}(\mathrm{d}\eta). \tag{4.2}$$

Since $\sum_{y=1}^{\mathtt{r}(\eta)} \Delta_\eta^k(y) \leqslant \mathtt{r}(\eta)$, we get $v_k \leqslant \frac{w_0}{\rho_k} < \infty$.

4.2 Equation for speeds from soliton interactions

We now prove that the speeds satisfy (1.7). We start by analyzing how collision affects the displacement of solitons. Let $m > k$, $\tilde{\gamma} \in \Gamma^m\eta$ and $\gamma \in \Gamma^k\eta$. We observe that $\tilde{\gamma}$ takes over γ in two time steps, both when $\tilde{\gamma}$ is free to move forward, i.e., when it is not nested inside a bigger soliton. 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 $\tilde{\gamma}$. During these two steps, γ is prevented from moving due to the collision with $\tilde{\gamma}$. For each step in between, γ is also prevented from moving, but this is due to collision between γ and the largest soliton in the nesting chain (the only which is free to move), not with $\tilde{\gamma}$ which is also prevented from moving. So collision with $\tilde{\gamma}$ causes γ to stay put at two time steps. Conversely, collision with γ causes $\tilde{\gamma}$ to move $2k$ units more than it normally would have.

Let $\gamma \in \Gamma_k\eta$ be such that $0 = x(\gamma)$. After t steps we have

$$x(\gamma^t) = k\,t - 2k\sum_{m>k} N_t^m + \sum_{m<k} 2m\,N_t^m. \tag{4.3}$$

Here N_t^m counts the number of m-solitons which collide with the tagged soliton γ during steps $1, \dots, t$. The first and last of such collisions is counted as a half in case only one of the two steps described above occurs in this lapse, so N_t^m may be a half integer.

Assume $\bar{\rho}_m = 0$ for all large m. By (1.6) and (4.3), to get (1.7) it suffices to show that

$$\tfrac{1}{t}N_t^m \to \bar{\rho}_m|v_k - v_m| \quad \text{ in probability under } \widehat{\mu}_k. \tag{4.4}$$

Let $m > k$. Consider the rightmost m-soliton of η contained in $(-\infty, -1]$ and let X_t^- be its position at time t. Likewise, consider the leftmost m-soliton of η contained in $[1, +\infty)$ and let X_t^+ be its position at time t. Now N_t^m is bounded from above by the number of m-solitons in $T^t\eta$ which intersect $[x(\gamma^t), X_t^+]$ and bounded from below by the number of m-solitons which are contained in $[x(\gamma^t), X_t^-]$. By (1.6), $\frac{1}{t}x(\gamma^t) \to v_k$ and $\frac{1}{t}X_t^\pm \to v_m$, $\widehat{\mu}_k$-a.s. Hence, by definition of

$\bar\rho$ and the Ergodic Theorem, the upper and lower bounds are both close to $\bar\rho_m(v_m - v_k)$, with high probability as t increases. This proves (4.4) for $m > k$. The case $m < k$ is analogous.

To complete the proof of (1.7), it remains to drop the assumption that $\bar\rho_m = 0$ for all large m.

The above proof contains all the argument, except that combining (4.3) and (4.4) requires a limit and an infinite sum to commute. To prove the general case, we take expectation in (4.3) with $t = 1$ and combine it with (4.1):

$$\begin{aligned} v_k &= \int x(\gamma^1)\widehat\mu_k(\mathrm{d}\eta) = \int \Big[k - \sum_{m>k} 2k\, N_1^m + \sum_{m<k} 2m\, N_1^m\Big]\, \widehat\mu_k(\mathrm{d}\eta) \\ &= k + \sum_{m\neq k} c_{k,m} \int N_1^m(\eta)\, \widehat\mu_k(\mathrm{d}\eta), \end{aligned}$$

where $c_{k,m} = -2k$ for $m > k$ and $2m$ for $m < k$. So it remains to show that $\int N_1^m(\eta)\,\widehat\mu_k(\mathrm{d}\eta) = \bar\rho_m |v_k - v_m|$. Similar to the argument in §4.1, we decompose

$$N_{t+1}^m(\eta) = N_1^m(\eta) + N_1^m(\widehat T_k \eta) + N_1^m(\widehat T_k^2 \eta) + \cdots + N_1^m(\widehat T_k^t \eta).$$

By $\widehat T_k$-invariance, $\frac{1}{t} N_t^m(\eta)$ converges $\widehat\mu_k$-a.s. to a random variable (i.e. a measurable function of η) whose average is $\int N_1^m(\eta)\,\widehat\mu_k(\mathrm{d}\eta)$. Now this variable is in fact constant and equal to $\bar\rho_m|v_k - v_m|$ by (4.4), concluding the proof.

4.3 Recursion formulas for independent components

We finally prove (1.8). 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, $\widehat\mu$ is the law of $\eta = M^{-1}\zeta$.

Note that (3.9) and (3.11) give the first two equations in (1.8). The third equation can be taken as the definition of s_k, it is a simple recursive definition once one has ρ, w and α. Combining these with (4.2), to get the last equation in (1.8) we need to show that

$$E \sum_{y=1}^{\mathtt{r}(\eta)} \Delta_\eta^k(y) = \alpha_k \cdot s_k. \tag{4.5}$$

We now use the assumption that $\eta = M^{-1}\zeta$, where M^{-1} denotes de reconstruction map of §2.3. First note that $\sum_{y=1}^{\mathtt{r}(\eta)} \Delta_\eta^k(y)$ equals the sum over all k-solitons γ in the first excursion of η of the size of the jump of γ in one time step, that is $x(\gamma^1) - x(\gamma)$. Moreover, k-solitons appended to k-slots belonging to bigger solitons will stay frozen, just switching zeros for ones, and only the k-solitons which are appended directly to the 0-th k-slot at $x = 0$ will actually jump. Furthermore, the size of their jump equals the distance between their leftmost one and their leftmost zero, 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 0-th k-slot is exactly $\zeta_k(0)$, which on average equals α_k by (3.10). So to conclude the proof of (4.5), it is enough to observe that s_k given by the recursion relation in (1.8) in fact gives the average distance between the tips of the head and tail of a typical k-soliton.

We show a stronger statement: that s_k equals both the average distance between the tip of the head and the tip of the tail as well as the average distance between the tip of the tail and the next k-slot. We made the statement stronger so we can prove it by induction. 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. 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 which are determined by $\{\zeta_m\}_{m>1}$. 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 (1.8).

4.4 Vertical speed

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 (2.3), we define displacement of a tagged k-slot $\pi \in S_k\eta$ measured in terms of records by

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

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 π^t in $T^t\eta$, so y_k^t also measures the displacement of tagged k-solitons.

Theorem 4.1. *Let μ be a measure on $\mathcal{X}$ such that under $\widehat{\mu}$ each k-component $M_k\eta$ is i.i.d. and they are independent over k. There exists a non-decreasing deterministic sequence $h = (h_k)_{k\geqslant 1}$ such that, μ-a.s. on η, for all $k \in \mathbb{N}$ and $\pi \in S_k\eta$,*

$$\lim_{t\to\infty} \frac{y_k^t(\eta, \pi)}{t} = h_k \in [k, \infty]. \tag{4.6}$$

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

$$h_k = k + \sum_{m>k} 2(m-k)(h_m - h_k)\rho_m, \quad k \geqslant 1. \tag{4.7}$$

The asymptotic speed of tagged records is given by

$$\lim_{t\to\infty} -\frac{\beta^t}{t} = v_0 := \sum_{m\geqslant 1} 2m\rho_m h_m, \quad \textit{for all } \beta \in R\eta,\ \mu\textit{-a.s.} \tag{4.8}$$

The asymptotic speed v_k of tagged k-slots is also given by

$$\lim_{t\to\infty} \frac{\pi^{k,t}}{t} = v_k = h_k w_0 - v_0, \quad \textit{for all } \pi \in S_k\eta,\ \mu\textit{-a.s.}, \tag{4.9}$$

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

$$v_0 = \sum_{m\geqslant 1} 2m\rho_m v_m. \tag{4.10}$$

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

$$\lim_{t\to\infty} -\frac{T^t\xi(0)}{t} = h_0 := \frac{v_0}{w_0}. \tag{4.11}$$

We outline the proof of the first four equations, referring to arXiv:1806.02798v3 for the details.

The proof of (4.6) is similar to that of (1.6). Each time an m-soliton overtakes a k-soliton, this causes the position of the k-soliton measured in records to be incremented by an extra factor of $2(m-k)$. On the other hand, the position of a soliton measured in records is not affected by overtaking smaller solitons. These two facts explain the origin of (4.7).

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 (4.6). Summing over m we get (4.8).

Finally, by (4.6), the tagged k-slot $\pi = 0 \in S_k\xi$ 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 (4.8) 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 (4.9).

From (4.9) we have $h_m = \frac{v_0+v_m}{w_0}$, substituting into (4.8) and using (3.12) we get (4.10). To prove (4.11) 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.

5 Postponed proofs

Proof of Proposition 3.1. By (2.1), this statement corresponds to the discrete version of Pitman's $2M - X$ theorem, see [HMO01, Pit75]. We will prove it here following the proof of Burke's Theorem given in [Rei57]. Recall the nearest-neighbor walk $\xi = \xi[\eta]$, which is in this case distributed as a simple random walk with i.i.d. increments distributed as Bernoulli(λ). We introduce the reflected process: $W(x) = \xi(x) - \min_{y\leqslant x} \xi(y)$, $x \in \mathbb{Z}$. Then for all x, $W(x)$ has distribution Geom$(\frac{\lambda}{1-\lambda})$, that is, $\mathbb{P}(W(x) \geqslant k) = (\frac{\lambda}{1-\lambda})^k$. Moreover, W is *reversible*. Observe that we can recover η from W: $A := \{x \in \mathbb{Z} : \eta(x) = 1\} = \{x \in \mathbb{Z} : W(x) - W(x-1) = 1\}$. We also write $A^c := \{x \in \mathbb{Z} : \eta(x) = 0\} := D \cup R$, where $D := \{x \in \mathbb{Z} : W(x) - W(x-1) = -1\}$ and $R := \{x \in \mathbb{Z} : x \text{ is a record}\} = \{x \in \mathbb{Z} : W(x) = 0\}$. Notice that $T\eta$ consists in reversing the increments of the boxes in A and D; in other words, exchange the roles of A and D. But the reversibility of W implies that A has the same distribution as D. The conclusion then follows. □

Proof of Proposition 1.4. 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.

> Start with the word $\ominus^\infty$ which is semi-infinite to the left
> **for** each symbol in the finite configuration η **do**
> | Append the symbol to the word
> | Perform annihilation if the two last runs have the same length
> | Symbols that annihilate correspond to a soliton
> **end**

For example, for the finite sequence $\eta = \oplus\oplus\ominus\oplus\oplus\ominus\ominus\oplus\oplus\oplus\ominus\ominus\ominus\ominus\ominus$ the algorithm would produce the words $\ominus^\infty\oplus$, $\ominus^\infty\oplus^2$, $\ominus^\infty\oplus^2\ominus$, $\ominus^\infty\oplus^2\underline{\ominus\oplus}$, $\ominus^\infty\oplus^3$, $\ominus^\infty\oplus^3\ominus$, $\ominus^\infty\oplus^3\ominus^2$, $\ominus^\infty\oplus^3\ominus^2\oplus$, $\ominus^\infty\oplus^3\underline{\ominus^2\oplus^2}$, $\ominus^\infty\oplus^4$, $\ominus^\infty\oplus^4\ominus$, $\ominus^\infty\oplus^4\ominus^2$, $\ominus^\infty\oplus^4\ominus^3$, $\ominus^\infty\underline{\oplus^4\ominus^4}$, and $\ominus^\infty\ominus = \ominus^\infty$, identifying a 1-soliton, a 2-soliton and a 4-soliton. For the example in Fig. 1.1, it produces $\ominus^\infty\oplus$, $\ominus^\infty\oplus^2$, $\ominus^\infty\oplus^3$, $\ominus^\infty\oplus^4$, $\ominus^\infty\oplus^4\ominus$, $\ominus^\infty\oplus^4\ominus^2$, $\ominus^\infty\oplus^4\ominus^2\oplus$, $\ominus^\infty\oplus^4\ominus^2\underline{\oplus\ominus}$, $\ominus^\infty\oplus^4\ominus^2\oplus$, $\ominus^\infty\oplus^4\underline{\ominus^2\oplus^2}$, $\ominus^\infty\oplus^5$, $\ominus^\infty\oplus^5\ominus$, $\ominus^\infty\oplus^5\underline{\ominus\oplus}$, $\ominus^\infty\oplus^5\ominus$, $\ominus^\infty\oplus^5\ominus^2$, $\ominus^\infty\oplus^5\ominus^3$, $\ominus^\infty\oplus^5\ominus^4$, $\ominus^\infty\oplus^5\ominus^4\oplus$, $\ominus^\infty\oplus^5\ominus^4\oplus^2$, $\ominus^\infty\oplus^5\ominus^4\oplus^2\ominus$, $\ominus^\infty\oplus^5\ominus^4\oplus^2\underline{\ominus\oplus}$, $\ominus^\infty\oplus^5\ominus^4\oplus^3$, $\ominus^\infty\oplus^5\ominus^4\oplus^3\ominus$, $\ominus^\infty\oplus^5\ominus^4\oplus^3\ominus^2$, $\ominus^\infty\oplus^5\ominus^4\underline{\oplus^3\ominus^3}$, $\ominus^\infty\underline{\oplus^5\ominus^5}$, identifying three 1-solitons, a 2-soliton, a 3-soliton, and a 5-soliton.

Let us call $\oplus$-*alternating suffix* (or simply $\oplus$-*suffix*) a finite word ω which is either empty or starts with $\oplus$ and such that each run in the word is strictly longer than the next one. So the above algorithm always produces words given by $\ominus^\infty$ followed by a $\oplus$-suffix. We define $\ominus$-*suffix* in the obvious way. The *net value* $v(\omega)$ of a finite suffix ω is the number of $\oplus$'s minus the number of $\ominus$'s.

Remark 1. The net value of a non-empty $\oplus$-suffix ω is positive and it is at most equal to the length $\ell_1(\omega)$ of its first run (e.g. for $\cdots\oplus^4\ominus^3\oplus$ we have $0 < 2 \leqslant 4$). In particular, $v(\omega) = \ell_1(\omega)$ only if it consists of a single run.

Remark 2. The net value of a finite suffix ω equals the net value of the portion of η that generated it, which in turn is given by the net increase in $\xi[\eta]$.

Remark 3. If the suffixes $\omega_1, \dots, \omega_n$ produced while processing a certain piece of η are all $\oplus$-suffixes, then $\ell_1(\omega_n)$ equals the maximal net value of ω_i for $i = 1, \dots, n$. In particular, if $v(\omega_n) = \max_i v(\omega_i)$, then $\ell_1(\omega_n) = v(\omega_n)$ and, by Remark 1, ω_n consists of a single run.

To prove the proposition we will split a finite η into three blocks and analyze how they interact under the data stream algorithm, both before and after the application of T, as shown in Fig. 5.1.

Define the first non-empty soft excursion as the piece of η going from the first $\oplus$ until the first point that makes the net value equal zero. Split this excursion into *rising* and *falling* parts as follows. The rising part goes until the point where the net value k is maximal (in case the maximum is attained more than once, take the rightmost one), and the falling part consist of the remaining boxes, until the end of the first soft excursion. The *remainder* consists of all the sites to the right of the falling block. Let $I_1, I_2, I_3 \subseteq \mathbb{Z}$ denote these sets of sites.

By definition of I_1 and by Remark 2, the streaming algorithm applied to η on I_1 always produces a non-empty $\oplus$-suffix, its net value is always at most k and ends being equal to k.

By Remarks 1 and 3, the word produced by the algorithm after processing this first block is $\oplus^k$. By similar considerations, the algorithm applied to η on I_2 always produces non-empty $\ominus$-suffixes whose net values are strictly between $-k$ and 0, except for the final step when it produces $\ominus^k$.

Hence, when processing η on $I_1 \cup I_2$, the $\oplus^k$ obtained after processing the rising part is kept untouched until the very end, when it is annihilated by the $\ominus^k$ obtained after processing the falling part. So when the algorithm starts processing η on I_3 there is no suffix left by the previous steps and this part of η is decomposed into solitons just as it would if it was processing $\eta_{|I_3}$ instead.

Now notice that, by the definition of T on $\xi[\eta]$, the net value of $T\eta$ on any prefix of I_3 is non-positive. Indeed, at the rightmost site y of I_2, the walk ξ coincides with its running minimum, so $T\xi(y) = \xi(y)$ and $T\xi(x) \leqslant T\xi(y)$ for all $x > y$. Hence, applying the streaming algorithm to this portion of $T\eta$ produces a $\ominus$-suffix at all steps.

Also, since $\xi(x) \geqslant \xi(y)$ for all $x \in I_1 \cup I_2$, by definition of T we have that η and $T\eta$ are the complementary 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 $T\eta$

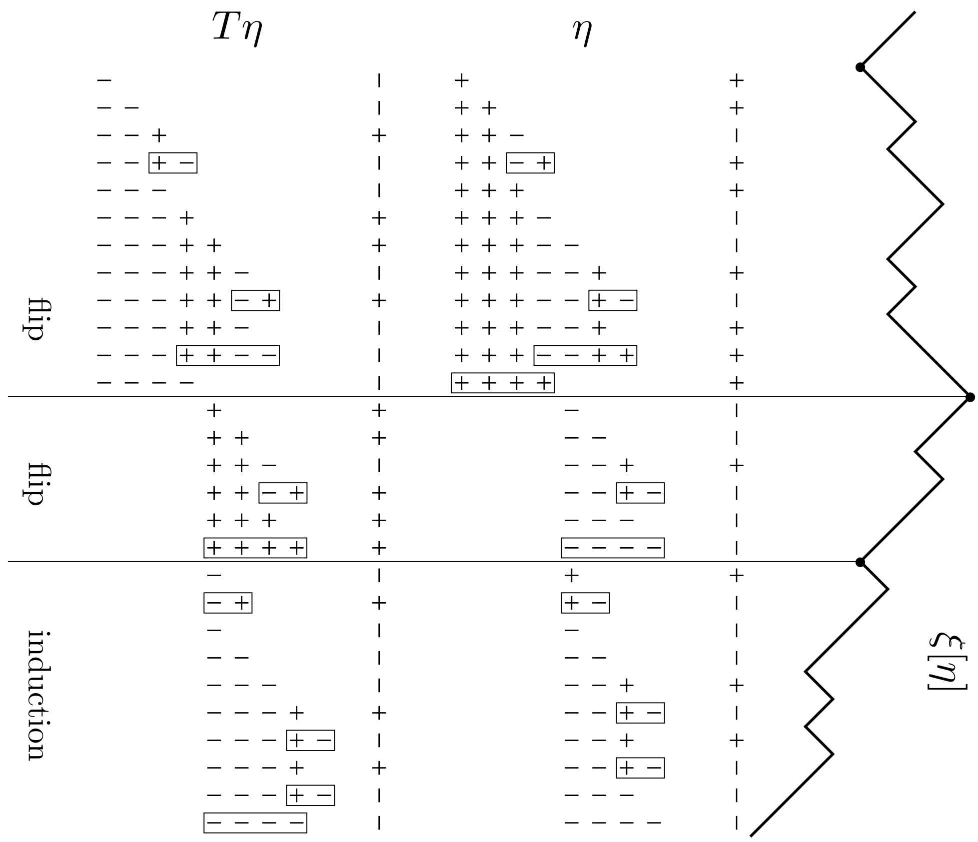

Figure 5.1: Example showing conservation of solitons by splitting space in three parts: rising, falling and remainder. After applying T, the configuration on the rising and falling parts are flipped, the smaller solitons are conserved and flipped, the biggest soliton moves forward, and will have its tail in the remainder part. Applying T to the remainder part conserves solitons by induction.

on I_1 is incorporated into the infinite prefix $\ominus^\infty$, and it will not annihilate with 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 $\ominus^\infty$. Hence, the presence of this $\oplus^k$ does not change how the algorithm processes $T\eta$ on I_3, neither before nor 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 all 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 $\eta'(x) := \eta(x)\mathbb{1}_{[y_1,y_4]}(x)$ Since 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 $[y_1, y_4]$. Since $\eta' \leqslant \eta$, and x being a record for η is a non-decreasing property in 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 $[y_2, y_3]$, $T\eta' \leqslant T\eta$ 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 $\gamma \in \Gamma_k\eta'$ 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 $[y_2, y_3]$. Since η' is a finite configuration, by the previous case this concludes the proof. □

Acknowledgments

We thank Roberto Fernández for very helpful discussions.

References

[CKST18] D. A. Croydon, T. Kato, M. Sasada, S. Tsujimoto. *Dynamics of the box-ball system with random initial conditions via Pitman's transformation*, 2018. Preprint. arXiv:1806.02147.

[FG18] P. A. Ferrari, D. Gabrielli. *Invariant measures for BBS with independent soliton components*, 2018. In preparation.

[Har71] T. E. Harris. *Random measures and motions of point processes.* Z. Wahrscheinlichkeitstheorie verw. Geb. **18**:85–115, 1971. doi.

[HMO01] B. Hambly, J. B. Martin, N. O'Connell. *Pitman's $2M - X$ theorem for skip-free random walks with Markovian increments.* Electron. Commun. Probab. **6**:73–77, 2001. doi.

[IKO04] R. Inoue, A. Kuniba, M. Okado. *A quantization of box-ball systems.* Rev. Math. Phys. **16**:1227–1258, 2004. doi.

[IKT12] R. Inoue, A. Kuniba, T. Takagi. *Integrable structure of box-ball systems: crystal, Bethe ansatz, ultradiscretization and tropical geometry.* J. Phys. A **45**:073001, 64, 2012. doi.

[KTZ17] T. Kato, S. Tsujimoto, A. Zuk. *Spectral analysis of transition operators, automata groups and translation in BBS.* Commun. Math. Phys. **350**:205–229, 2017. doi.

[LLP17] L. Levine, H. Lyu, J. Pike. *Phase transition in a random soliton cellular automaton*, 2017. Preprint. arXiv:1706.05621.

[LPS14] T. Lam, P. Pylyavsky, R. Sakamoto. *Rigged configurations and cylindric loop Schur functions.* Ann. Inst. Henri Poincaré Comb. Phys. Interact. **(to appear)**, 2014. arXiv:1410.4455.

[MIT06] J. Mada, M. Idzumi, 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**:053507, 18, 2006. doi.

[Pit75] J. W. Pitman. *One-dimensional Brownian motion and the three-dimensional Bessel process.* Advances in Appl. Probability **7**:511–526, 1975. doi.

[PS73] S. C. Port, C. J. Stone. *Infinite particle systems.* Trans. Amer. Math. Soc. **178**:307–340, 1973. doi.

[Rei57] E. Reich. *Waiting times when queues are in tandem.* Ann. Math. Statist. **28**:768–773, 1957. doi.

[Sak14a] R. Sakamoto. *Rigged configurations and Kashiwara operators.* Symmetry Integr. Geom. **10**:Paper 028, 88, 2014. doi.

[Sak14b] ———. *Ultradiscrete soliton systems and combinatorial representation theory.* RIMS Kokyuroku **1913**:141–158, 2014. url.

[Tho00] H. Thorisson. *Coupling, stationarity, and regeneration.* Probability and its Applications (New York). Springer-Verlag, New York, 2000. doi.

[TM97] D. Takahashi, J. Matsukidaira. *Box and ball system with a carrier and ultradiscrete modified KdV equation.* J. Phys. A **30**:L733–L739, 1997. doi.

[TS90] D. Takahashi, J. Satsuma. *A soliton cellular automaton.* J. Phys. Soc. Japan **59**:3514–3519, 1990. doi.

[TTMS96] T. Tokihiro, D. Takahashi, J. Matsukidaira, J. Satsuma. *From soliton equations to integrable cellular automata through a limiting procedure.* Phys. Rev. Lett **76**:3247, 1996. doi.

[TTS96] M. Torii, D. Takahashi, J. Satsuma. *Combinatorial representation of invariants of a soliton cellular automaton.* Physica D: Nonlinear Phenomena **92**:209 – 220, 1996. doi.

Pablo A. Ferrari[1], Chi Nguyen[1], Leonardo T. Rolla[12], Minmin Wang[1]

[1] Argentinian National Research Council at the University of Buenos Aires

[2] Corresponding Author. e-mail: leorolla@dm.uba.ar