THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Stochastic systems with locally defined dynamics

Anton Muratov

Department of Mathematical Sciences
Division of Mathematical Statistics
Chalmers University of Technology and University of Gothenburg
Göteborg, Sweden 2014

Stochastic systems with locally defined dynamics
Anton Muratov
ISBN 978-91-7597-040-0

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 3721
ISSN 0346-718X

Department of Mathematical Sciences
Division of Mathematical Statistics
Chalmers University of Technology and University of Gothenburg
SE-412 96 Göteborg
Sweden
Telephone +46 (0)31 772 1000

Typeset with LaTeX.
Printed in Göteborg, Sweden 2014

Stochastic systems with locally defined dynamics
Anton Muratov

Abstract

We study three different classes of models of stochastic systems with locally defined dynamics. Our main points of interest are the limiting properties and convergence in these models.

The first class is the locally interactive sequential adsorption, or LISA, models. We provide the general LISA framework, show that several classes of well-understood models fall within the framework, such as Polya urn schemes and fragmentation processes. We study several particular new examples of LISA processes having the feature of scalability. We provide the sufficient conditions for the existence of limiting empirical measures, and prove a bound for the speed of convergence.

The second class is Bit Flipping models, where we study a behaviour of a sequence of independent bits, each flipping between several states at a given rate p_k. We define two particular models, Binary Flipping and Damaged Bits, and find the conditions on the rates $\{p_k\}$ at which the models switch from the transient to the recurrent behaviour; as well as provide bounds for moments of the recurrence time under a certain set of conditions in the recurrent case, and prove the central limit theorem.

The third class is Random Exchange Models where a countable collection of agents are trading independent random proportion of their masses with neighbours in a stepwise fashion. We find the stationary regimes for such models, and prove a limit theorem. As a corollary, we obtain a new invariance property of a stationary Poisson process on the real line with respect to a certain neighbour-dependent point shift.

Keywords: point process, sequential adsorption, stopping set, random measure, Polya urn, convergence of empirical measures, bit flipping, recurrence, renewal process, Poisson process, Dirichlet distribution, random matrices

Acknowledgment

First of all, I would like to thank my supervisor, Sergey Zuyev, for always being there for me. This thesis would be impossible without his help. I thank Sergey Foss for his support and strategic guidance.

I thank the staff at Chalmers Matematiska Vetenskaper for inspiring discussions, comments, and a number of excellent insightful courses which meant so much in my mathematical education. In particular, Olle Häggström, Jeff Steif, Serik Sagitov, Aila Särkkä, Torgny Lindvall, Mats Rudemo, Holger Rootzen, Peter Sjögren, thank you! I would like to thank Serik Sagitov and Petter Mostad for an opportunity to read lectures, which was always very demanding and motivating for me.

I would like to thank the administrative staff at MV, in particular Marianne, Fia, Lotta, Setta, Maria, Gittan and Cecilia, for their flawless management and making the administrative matters, that otherwise would have been cumbersome and impregnable, easy and straightforward.

I thank my friends and colleagues here at Matematiska Vetenskaper for being a wonderful company during these five years. Dmitrii, Alexandra, Sofia, Malin, Magnus, Peter, Vera, Krzysztof, José, Adam, Mehdi, Urban, Alexei, Dmitrii, Viktor, Frederik, Marianna, Henrike, Ivar, Malin, Claes, and many more, thank you, it was a great pleasure to share this time here with you.

Last but not least, I would like to thank my family, constantly supporting me from the beginning of this long road. I thank Natasa for her love.

Anton Muratov
Göteborg, June 18, 2014

List of Papers

The licentiate thesis includes the following papers.

I. *LISA: Locally Interacting Sequential Adsorption*, **Muratov, A.**, Zuyev, S., Stochastic Models Vol. 29, Iss. 4, (2013)

II. *Bit flipping models and time to recover*, **Muratov, A.**, Zuyev, S. (2014)

III. *Neighbour-dependent point shifts and random exchange models: invariance and attractors*, **Muratov, A.**, Zuyev, S. (2014)

Contents

1 Introduction

This thesis contains three papers, each devoted to a separate class of new models of stochastic processes. The main points of interest are the limiting behaviour of the models when time goes to infinity, the existence of the stationary regime and its properties.

Section 2 is a review of the paper on LISA processes. It starts with an original R. Darling construction, which gave us an idea to consider the LISA models. We introduce the general LISA setting, and the particular examples we consider, along with the brief overview of some other models falling into the setting. The topics include Polya's urn models and their connection to Dirichlet processes, branching random walks, unity-splitting process and some other fragmentation processes. In the end of the first section we outline our main results on the models considered.

Section 3 is a review of the second paper of the thesis on Bit Flipping models. We outline our original motivation, explaining how a Bit Flipping problem naturally arises in the analysis of a certain LISA-related problem. We indicate a link between bit flipping and the areas of dynamic percolation, random walks on an infinite cube, and analysis of the algorithms. We then proceed with formulating the two models, Binary Flipping (BF) and Damaged Bits (DB), and our results.

Section 4 contains some background for the third paper. We introduce motivation for the Random exchange models, make an introduction to Dirichlet random vectors and their properties. We briefly summarize our main results for Random exchange models, and their implications to the neighbour-dependent random shifts of stationary point processes on the real line.

2 LISA: Locally Interacting Sequential Adsorption

The main inspiration for studying LISA models was a particular problem proposed by Richard Darling to my scientific advisor Sergey Zuyev. The original model is defined as follows.

Let $X_0 = \{x_1, x_2, \ldots, x_{n_0}\} \subset \mathbb{R}^2$ be the starting configuration, where all of the points are disjoint. At each time step $n = n_0 + 1, n_0 + 2, \ldots$ we are adding a new point to a configuration according to the following algorithm:

- Sample a random index χ_n, uniformly distributed over the indices $1, 2, \ldots, n$. The point x_{χ_n} will generate a new point.
- Find the smallest ball $B(x_{\chi_n}, R)$, containing $m, m \geq 3$ points from $X_n \setminus x_{\chi_n}$.
- Assume that m closest neighbours of x_{χ_n} lying in $B(x_{\chi_n}, R)$ come from a normal distribution with mean x_{χ_n}. Estimate the covariance matrix C of that distribution from those m points.
- Sample a new point x_{n+1} from a normal distribution $\mathrm{N}(x_{\chi_n}, C)$ with mean x_{χ_n} and the estimated covariance matrix C.
- Add the new point to the configuration: $X_{n+1} = X_n \cup \{x_{n+1}\}$

The resulting model has a few interesting properties.

The dynamics of a system is defined locally: the location of each new point only depends on the local configuration around its parent. Therefore, we can, for example, embed the whole process in a continuous time, letting every particle generate new ones with fixed intensity, and distributing them around itself according to local rules. That reflects the possible real-world applications of a model: growth of the cities on the plane, bacteria growth etc.

As simulations show, see Figure 1,the geometry of a configuration positively reinforces itself: if we start from a "round-shaped" configuration, then the covariance matrices will be isotropic and hence the whole configuration will likely remain round-shaped. But every now and then the configuration shoots out "beams" of particles, in such "beam" areas the estimated covariance matrices are highly anisotropic, i.e. with high concentration along some vector, and that behaviour is also self-reinforced.

One more significant feature of the model is scaling invariance, which might imply stochastic self-similarity and various fractal properties for the limiting empirical measure.

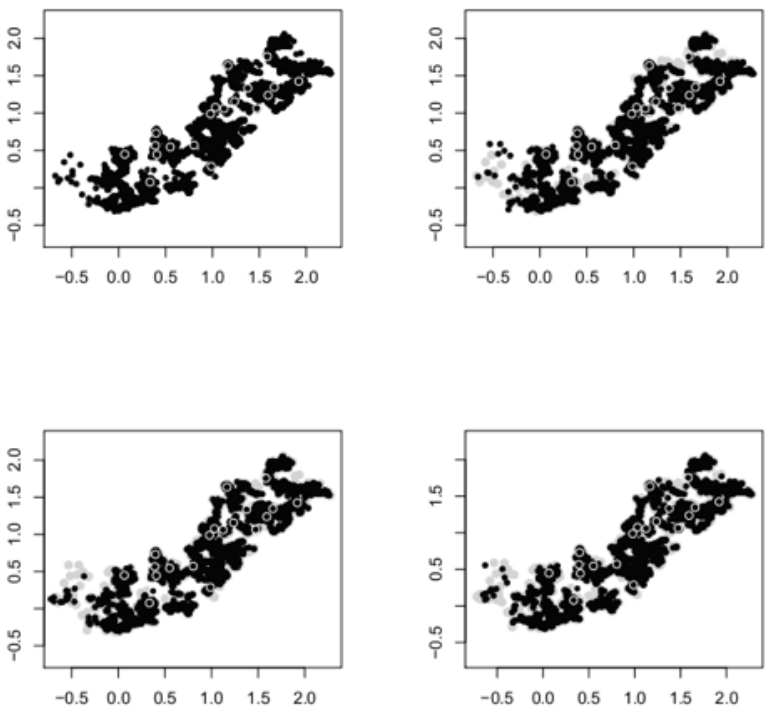

Figure 1: The sequence shows the simulated process after $n = 10^4, 2{\cdot}10^4, 3{\cdot}10^4, 4{\cdot}10^4$ steps. Newly added particles are shown in dark, previously existing – in light grey, initial particles are contoured void circles.

2.1 General LISA setting

We believe that the proper tool for handling the local dependence structure for special point processes is the notion of a stopping set. It extends the concept of a stopping time for Markov processes in one time-dimension. The definition is borrowed from [22].

Let W be a locally compact separable topological space, with $\mathbb{K}$ and $\mathbb{F}$ being its system of compact and closed sets, respectively. Let $(\Omega, \{\mathcal{F}_K\}, \mathbf{P})$ be a filtered probability space. Filtration $\{\mathcal{F}_K\}$ here is a collection of σ-algebras $\mathcal{F}_K$ indexed by compact sets $K \in \mathbb{K}$ that has the following properties:

- monotonicity: $\mathcal{F}_{K_1} \subseteq \mathcal{F}_{K_2}$ for any two compact $K_1 \subseteq K_2$;
- continuity from above: $\mathcal{F}_K = \cap_{n=1}^{\infty} \mathcal{F}_{K_n}$ if $K_n \downarrow K$.

A random closec set S is a measurable mapping $S{:}(\Omega, \{\mathcal{F}_K\}, \mathbf{P}) \mapsto [\mathbb{F}, \sigma_f]$, where σ_f is the σ-algebra generated by the system $\{F{\in}\mathbb{F} : F \cap K \neq \varnothing\}$, $K \in \mathbb{K}$.

Definition 2.1. A random closed set S is called a stopping set, if for any $K \in \mathbb{K}$ the event $\{S \subseteq K\}$ is $\mathcal{F}_K$-measurable.

More specifically, a random closed set S is a stopping set with respect to the sigma-algebra generated by a point process X if in order to know the shape of S one only needs to know the configuration of X inside of it. A simple example of a stopping set, illustrating this intuition, and the one we will be using the most, is the minimal closed ball centred in $x \in W$ containing at least k points of a point process.

A more involved example comes from Voronoi diagrams. If $X \subset \mathbb{R}^2$ is a locally finite point process, and $x \in \mathbb{R}^2$ is an arbitrary point then the Voronoi cell containing point x is not a stopping set, because in order to know its shape we need the information from outside the cell: in particular we need to know the configuration of X inside the shaded region on Figure 2. The shaded region is called Voronoi flower corresponding to the cell, and it is a stopping set.

Having the notion of stopping sets at hand, we can now define the general setting for LISA processes. Start with some initial configuration $X_{n_0} = \{x_1, x_2, \ldots, x_{n_0}\}$ of points in the sample space W. Typically, $W = \mathbb{R}^d$, but other options are also possible. The core of the process is the procedure of generation of the new points. We follow the general scheme:

- At time n, $n = n_0, n_0 + 1, n_0 + 2, \ldots$, pick a parent point $x \in X_n$ uniformly at random among currently existing.
- Find a stopping set $S_n = S(x, X_n)$ defined by the configuration at time step n and previously picked x.

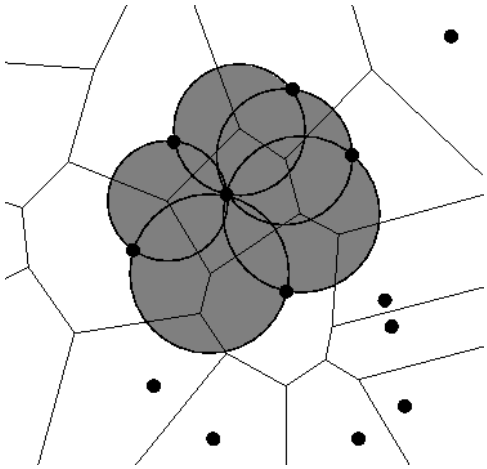

Figure 2: An example of a stopping set: Voronoi flower

- For a parent point x, sample a random variable $\zeta_x = \zeta(S_n)$, whose distribution is defined by the geometry of a stopping set and which is otherwise not dependent on X_n:

$$\mathbf{P}(\zeta_x \in B | X_n) = \mathbf{P}(\zeta_x \in B | S_n)$$

- Add the new point $x_{n+1} = \zeta_x$ to the configuration:

$$X_{n+1} = X_n \cup \{x_{n+1}\}$$

The natural idea for the continuous-time alternative, which tracks back to [1], is to embed the procedure of generation of the new points into a continuous-time branching process, so that if X_t denotes the configuration at time $t > 0$, every point $x \in X_t$ generates children at the unit rate. The location at which the child point is placed is then controlled as before by the geometry of its parent's stopping set $S(x, X_t)$ at the moment of generation.

As we see further, the above setting is quite flexible, in particular, we can obtain any kind of a stepwise growing Markov point process by just setting the stopping set to be the whole configuration: $S_n = X_n, n = n_0, n_0 + 1, \ldots$. Tweaking S_n and $\zeta(S_n)$, one could obtain variants of Polya's urn models, branching random walks, segmentation processes, etc.

In our paper we mainly focus our attention on the following particular examples of LISA models, possessing the properties of scale invariance.

Example 1 Put $W = [0,1] \subset \mathbb{R}$, $n_0 = 1$, and $X_{n_0} = \{0\}$. Denote by $r(x)$ the rightmost neighbour of the point x, or 1 if x does not have any neighbours to the right:

$$r(x) := \min\{y \in X_n \backslash \{x\} \cup 1 : y > x\},$$

and let the stopping set $S(x, X_n)$ be the interval from the point to its right neighbour: $S(x, X_n) = [x, r(x)]$. Put $\zeta(S) \sim \mathrm{Unif}(S)$, so that the new point is distributed uniformly on an interval to the right from the parent point.

Example 1 can be regarded as a stick-breaking process. It is directly related to the construction of Dubins and Freedman [8]: the distribution function of the limit of an empirical measure is exactly the random distribution function in [8]. This example in a slightly broader generality is thoroughly studied in [20]. One of the particularly interesting features of this model is the singularity of the resulting limiting measure with the explicitly given Hausdorff dimension of the support, as found in [20]. We use this example as an illustration of possible properties of LISA processes.

Example 2 Put $W = \mathbb{R}$, $n_0 \geq 2$, fix $X_{n_0} \subset \mathbb{R}$ consisting of n_0 distinct particles. Define the distance to the closest neighbour $d(x, X) = \min\limits_{y \in X \backslash \{x\}} |x - y|$, the stopping set $S(x, X) = [x - d(x, X), x + d(x, X)]$, and the distribution of a new point $\zeta_x \stackrel{\mathcal{D}}{=} x + d(x, X)\varepsilon$, where $\varepsilon = \pm 1$ with probability $1/2$, and ψ is a random variable concentrated on $(0, 1)$, so the newly generated particle stays inside the stopping set.

Example 3 A one-dimensional variation of the model described in the Introduction. Take W, X_{n_0}, $S(x, X)$ and $d(x, X)$ as in the previous example, and ζ_x Normally distributed with mean x and standard deviation $ad(x, X)$ for some constant parameter $a > 0$.

Example 4 Put $W = \mathbb{R}^d$, $n_0 \geq 2$, fix $X_{n_0} \subset \mathbb{R}^d$. As before, $d(x, X) = \min\limits_{y \in X \backslash \{x\}} ||x - y||$ is the Euclidean distance to the closest neighbour. The stopping set $S(x, X)$ is defined as a closed ball with the centre in x and radius $d(x, X)$. It is the minimal closed ball, containing at least one point of $X \setminus \{x\}$. Define $\zeta(S(x, X))$ to be distributed as $x + \psi d(x, X)$, where ψ is some $\mathbb{R}^d$-valued random variable.

The examples 3 and 4 are simplifications of the original R. Darling model, where the stopping set only contains one closest neighbour. We provide our results for these two examples further in the text, after a brief review of the related topics.

Example 5 Let W be a measurable space, μ a probability measure on W. Define $S(x, X)$ to be the whole X. Random variable ζ_x is x with probability $1 - 1/n$, where $n = \operatorname{card} X$, otherwise it is distributed according to the parameter measure μ. The LISA process defined this way corresponds to Blackwell-McQueen construction [4], a generalization of the Pólya urn scheme. The limiting empirical measure of this procedure is a Dirichlet process with parameter measure μ.

2.2 Polya's urn schemes

The simplest version of a Polya's urn model is defined as follows. Let the urn initially contain k black balls and l white balls. Let $n_0 = k + l$. At each step, draw a ball from the urn, look at its colour and return it back to the urn, along with one more ball of the same colour. Then the proportion of black balls in the urn tends to a random limit, which is distributed as $\mathrm{Beta}(k, l)$. In particular, if $k = l = 1$, the limit is uniform on $(0, 1)$ interval.

This version of Polya urn model is a "degenerate" case of LISA setting, obtained by taking the sample space W to be the set of two colours, the stopping set $S(x, X) = x$ and $\zeta_{S(x,X)}$ to be degenerate in x. It is also a border case of R. Darling model: let us say that instead of placing a new point at a normally distributed distance scaled with the distance to the closest neighbour(s), we place each new point at the location of its parent. Then the points of initial configuration can be regarded as a set of colours, and the starting configuration as an urn, containing one ball of each of the n_0 colors. The rest of dynamics stays the same.

This simple Polya's urn model allows for multiple generalization, including different amount of colours and various rules for reinforcement. One possible generalization is having the replacement controlled by a reinforcement matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

That means, when we pull out a black ball from the urn, we return it back along with a additional black and b white balls, whereas if the pulled ball's colour is white, we return it together with c black and d white balls. Then, depending on the relations between a, b, c, d there are different cases. For the details one should refer to [13, 6, 10, 15]. For the limit theorems see [11, 10, 3]. For a survey of the processes with reinforcement and further links, see [19]

See [1] for the embedding of the Polya's urns into a continuous-time multiclass birth and death processes, and the related results on limit behaviour.

One of the implications of Polya's urn model is a probabilistic insight on how to construct random measures in general, [4], along with a straightforward way to simulate the Dirichlet processes, a rich class of

priors, widely used in statistics due to their analytic tractability. The construction in [4] involves an infinite amount of colours with the limiting measure being a sample from Dirichlet process with a corresponding parameter measure on the space of colours.

2.3 Random Splitting

Example 1 above is a variant of a stick-breaking process. This is a sequential random splitting of $[0,1]$ interval where at each step the sub-interval picked among the existing uniformly at random, is split in two intervals in a random location uniformly distributed inside of the interval. As mentioned above, this process in a slightly more general setting is studied in detail in [20].

The splitting of the unity in [20] is a variation of Kakutani's splitting procedure [14], and formally defined as follows. Let $\{\chi_n\}_{n\geq 1}$ be a sequence of i.i.d. random variables such that χ_n is uniform over $\{0,1,2,\ldots,n\}$. Let $\{W_n\}_{n\geq 1}$ be a sequence of i.i.d. random variables with values in $(0,1)$. Define X_n in a following way: $X_1=\{0,1\}$ and if $\xi_0\leq\ldots\leq\xi_n$ are the elements of X_n, enumerated in the increasing order, then

$$X_{n+1}=X_n\cup\{\xi_{\chi_n}+W_n(\xi_{\chi_n+1}-\xi_{\chi_n})\}$$

The limit of the empiric measures is a random singular measure, almost surely supported by a (random) set of Hausdorff dimension D which is found explicitly in [20]. In the original Kakutani's procedure [14] the longest interval is always selected for splitting, when as in [20] the next interval to split is selected uniformly among those existing at the moment. The limit in [14] is then a deterministic uniform measure. In continuous time, the two selection procedures correspond to either giving all of the intervals equal splitting rates, or giving each interval the rate of splitting proportional to its length, respectively.

2.4 Main results

The construction proposed by R. Darling fits into LISA setting, but seems to be mathematically non-tractable, due to a very intricate character of dependency. Instead, our main points of interest in the paper are the characteristics of the limiting measure in the Examples 2,3,4. When does the configuration stay inside a compact region? Does the limiting measure always exist, and what are its properties? We start with almost sure boundedness in Example 2.

Theorem 2.1. *Denote*

$$m_n=\min\{x:x\in X_n\},$$
$$M_n=\max\{x:x\in X_n\}$$

for the LISA model in Example 2. Then almost surely

$$-\infty < \lim_{n\to\infty} m_n \le \lim_{n\to\infty} M_n < +\infty$$

Next, we provide a sufficient condition for boundedness in Examples 3, 4. Let us recall the rule by which new particles are added:

$$x_{n+1} = x_{\chi_n} + d(x_{\chi_n}, X_n)\psi_n,\ n = n_0, n_0+1, \dots$$

Denote by $\eta_n = ||\psi_n|| \stackrel{\mathcal{D}}{=} \eta$ the Euclidean norm of ψ_n, introduce $\hat{\eta} = \min(\eta, 1)$, and a function $\varphi(s) = \mathbf{E}\,\eta^s + \mathbf{E}\,\hat{\eta}^s$.

Theorem 2.2. *In Examples 3 and 4, whenever $\varphi(s) < 1$ for some $s : 0 < s \le 1$, the configuration stays bounded almost surely:*

$$\sup_n ||x_n|| < \infty a.s.$$

We then move on to the limits of the empirical measure. Let $\nu_n = \sum_{1\le k\le n} \delta_{x_k}$ denote the empirical measure on the nth step, and μ_n denote the distribution of x_n – a new point added on the nth step. Introduce $D_{j,N,\infty}$ – a set of all points generated by point x_j after time N.

Theorem 2.3. *The condition*

$$\max_{1\le j\le N} \sup_{i\in D_{j,N,\infty}} ||x_i - x_j|| \to 0,\ N\to\infty$$

is sufficient for the almost sure existence of a weak limit measure μ^* for the sequence μ_n corresponding to LISA processes from Examples 2, 3, and 4. In this case, with probability 1,

$$\mu_n \Rightarrow \mu^*, \nu_n \Rightarrow \mu^*.$$

Corollary 2.1. *The limiting measure in Example 2 exists.*

Corollary 2.2. *The condition*

$$\varphi(s) < 1 \text{ for some } s \in (0, 1]$$

is sufficient for existence of the limiting measure in Examples 3,4.

We also prove a bound on the speed of convergence to the limiting measure, estimating the maximal spacing d_n^* of the configuration. The maximal spacing d_n^* is the maximal distance from some point $x \in X_n$ to its closest neighbour in X_n:

$$d_n^* = \max_{x\in X_n} d(x, X_n)$$

Introduce $\sigma = \sup_{s>0,\varphi(s)>0} \frac{1-\varphi(s)}{s}$.

Theorem 2.4. *In Examples 2, 3, 4, if $\varphi(s) < 1$ for some positive s, then*

$$\limsup_n n^\sigma d_n^* < \infty$$

3 Bit Flipping Models

3.1 Model description

The second paper of the thesis addresses the questions of limit behaviour of Bit Flipping models. This section briefly summarizes the original motivation and the main results.

We define the Binary Flipping (BF) and Damaged Bits (DB) models in the discrete time setting, and introduce the natural continuous time embedding. In both models we are faced with an infinite sequence of elements, or "bits", indexed with positive integers, where every bit can be in several possible states. In both models, we start with a sequence of bits in an idle state. At each step n, the index χ_n of the next bit to change its state is sampled independently of everything from a fixed distribution $\mathcal{P} = \{p_1 > p_2 > p_3 > \dots\}$ on positive integers. $\mathcal{P}$ is the only parameter in the model.

Binary Flipping (BF) In the first model we let the bits switch between the idle and active states, denoted by 0 and 1, respectively. Formally, let $\{\zeta_n\}$ be the Markov chain on the state space

$$\mathcal{X} = \{x \in \{0,1\}^{\mathbb{N}} : x \text{ has finitely many non-zeros}\}$$

where ζ_n^k denote the state of kth bit at time n for $k = 1, 2, \dots$, $n = 0, 1, 2, \dots$, and $\zeta_n^k = 0, k = 1, 2, \dots$, so all of the bits start in the idle state. Let $\mathcal{P} = \{p_1, p_2, \dots\}$ be the probability distribution on positive integers. At each time step $n = 1, 2, \dots$ sample the index χ_n from $\mathcal{P}$ independently of everything. The bit with the index χ_n is flipped, i.e. its state is changed to the opposite:

$$\textit{idle(0)} \leftrightarrow \textit{active(1)}$$

Formally, for $n = 0, 1, \dots$ and $k = 1, 2, \dots$

$$\zeta_{n+1}^k = \begin{cases} 1 - \zeta_n^k, & k = \chi_{n+1}, \\ \zeta_n^k, & k \neq \chi_{n+1}. \end{cases} \tag{1}$$

The main object of interest is the stopping time τ_{BF}, the time of return to the ground state where all of the bits are idle:

$$\begin{aligned} \tau_{\mathrm{BF}} &= \min\{n \geq 1 : \text{no bits are active at time } n\} \\ &= \min\{n \geq 1 : \zeta_n^k = 0, k = 1, 2, \dots\} \end{aligned}$$

Depending on τ_{BF} being a.s. finite or infinite with positive probability, we refer to BF model as recurrent or transient. This corresponds to recurrence or transience of Markov chain $\{\zeta_n\}$.

Damaged Bits (DB) In the second model we allow for three possible states: idle, active and damaged, corresponding to 0, 1 and 2. Now $\{\zeta_n\}$ is a Markov chain on the state space

$$\mathcal{Y} = \{y \in \{0,1,2\}^{\mathbb{N}} : y \text{ has only a finite number of non-zeros}\}$$

Again, let ζ_k^n denote the state of kth bit at time n and $\zeta_k^n = 0$. At time step n state of the bit with the index $\xi_n \sim \mathcal{P}$ is changed according to the dynamics

$$\textit{idle(0)} \to \textit{active(1)} \to \textit{damaged(2)}$$

Formally, for $n = 0, 1, \dots$ and $k = 1, 2, \dots$

$$\zeta_{n+1}^k = \begin{cases} \min\{\zeta_n^k + 1, 2\}, & k = \chi_{n+1}, \\ \zeta_n^k, & k \neq \chi_{n+1}. \end{cases}$$

The stopping time τ_{DB} we are interested in is the first time $n \geq 1$ of returning to the ground state when none of the bits are active:

$$\begin{aligned} \tau_{\mathrm{DB}} &= \min\{n \geq 1 : \text{ none of the bits are active}\} \\ &= \min\{n \geq 1 : \zeta_n^k \neq 1, k = 1, 2, \dots\} \end{aligned}$$

Again, depending on τ_{DB} being a.s. finite, or infinite with positive probability, we call DB model recurrent or transient. Here it does not correspond to recurrence or transience of Markov chain $\{\zeta_n\}$, because the ground state does not correspond to any one particular state of the Markov chain, rather it is a countable collection of states.

Continuous time representation As mentioned above, both BF and DB models admit natural continuous time implementations. With a slight abuse of notation, we can consider a vector $(\zeta_t^1, \zeta_t^2, \dots)$ where the coordinates are independent continuous time Markov jump processes, each with the state space {idle, active} (BF) or {idle, active, damaged} (DB). For each $k = 1, 2, \dots$ we let ζ_t^k start in the idle state and change its state with rate p_k, with the respective dynamics. Then if we denote by $\{t_n\}_{n \geq 0}$, the sequence of transition times of the compound Markov jump process $\{\zeta_t\}_{t \geq 0} = \{(\zeta_t^1, \zeta_t^2, \dots)\}_{t \geq 0}$ with $t_0 = 0$, the embedded Markov chain $\{\zeta_{t_n}\}$ is a distributional copy of the original BF or DB model, respectively. Further we use both discrete- and continuous-time setting to formulate our results, depending on which of the two is more convenient at the moment: it is particularly convenient, because the notions of recurrence/transience stay the same for both implementations.

3.2 Motivation: from LISA to flipping bits

The original model considered was born in the discussions with Sergey Foss while looking for an example of LISA model where the configuration grows in size indefinitely.

In LISA setting, let $W = \{0, 1, 2, \dots\}$, $X_{n_0} = X_1 = \{0\}$, let η be the non-negative integer-valued random variable with distribution $\mathcal{P} = \{p_1, p_2, \dots\}$, let $\psi_{S(x,X)}$ be distributed as η, if x is a single point in X and be degenerate in 0 if x is a multiple point. Here we allow multiple points, so the same point can present in the configuration in many instances. The configuration's support is growing as long as there exist single points, but as soon as every existing point is multiple, the new ones are only born in 0. The question of interest is, does every existing point almost surely becomes multiple at some moment of time, or is it possible for a sequence of single points to "escape" to infinity with possible probability? The answer depends on the asymptotic of the decay of distribution $\mathcal{P}$, as given by the analysis of the Damaged Bits model below.

The term "Bit Flipping" term is adopted from [2] where it is used in the context of analysis the behaviour of a random-edge simplex algorithm on a Klee-Minty cube. There, in a sequence of bits indexed with positive integers $\{1, 2, 3, \dots\}$ each bit is flipping its state at unit rate, and when one bit changes its state, all of the bits to the right from it do so as well. Note that although there is an infinite amount of events happening during any finite time interval, any finite consecutive subset of bits starting with the first is still a properly defined Markov jump process in a continuous time with a finite state space. Our models differ from the one in [2]: we do not have dependency structure, instead, we introduce the different rates of flipping.

Bit Flipping models are related to the dynamical percolation processes, see [12]. We can consider the bits in BF model as edges of the $\mathbb{Z}^1$ graph flipping between open and closed states at different rates. In this context, the question of recurrence/transience corresponds to the question of existence of an infinite sequence of percolation times on $\mathbb{Z}^1$, at which zero is connected to infinity.

Another possible application is in the reliability theory. Consider the following problem: assume we have a composite system with a big (or infinite) number of independent components, indexed with positive integers, where each of the components is working continuously, until it breaks; and that for every $k \in \mathbb{Z}$, kth component is broken after exponential time with rate p_k. When one of the components fails, the whole system stops working correctly. It seems reasonable to have each component checked for failure and replaced with the rate q_k proportional to the rate of failure. In that case, every components' state changes are described by an independent copy of the same Markov jump process, but with the time scale stretched proportionally to $1/p_k$. Binary Flipping

model is a special case of this model in a situation where the rates of failure and replacement are the same, $q_k = p_k$.

3.3 Main results

Our main points of interest in the Bit Flipping models are those related to the probability of the return to the ground state. It is straightforward that the above probability is greater than zero, but is it always one? If it is, is the expectation of the recurrence time finite? Can we give some estimate for the moments of the time of return? How does the number of active bits grow as time goes to infinity?

The model formulation is very simple, and yet the answers to most of the above questions turn out to be non-trivial. First, we address the question of recurrence. It turns out, that both BF and DB models exhibit recurrent or transient behaviour, depending on the speed of the decay of p_k.

Theorem 3.1. *If the distribution $\mathcal{P} = \{p_1, p_2, \dots\}$ is such that:*

(i) $\limsup_{k\to\infty} 2^k p_k < \infty$, *then BF model is recurrent, i.e.* $\tau_{\mathrm{BF}} < \infty$ *a.s.*

(ii) $\liminf_{k\to\infty} (2-\varepsilon)^k p_k > 0$ *for some* $\varepsilon > 0$, *then BF model is transient, i.e.* $\mathbf{P}\{\tau = \infty\} > 0$.

Loosely speaking, Theorem 3.1 means that the critical decay of BF model is Geometric distribution with parameter $p = 0.5$. A non-random version of the model provides an illustration of the possible mechanics behind this.

Assume we have an infinite sequence of bits in state 0, indexed with $\{1, 2, \dots\}$. Let kth bit flip its state every 2^{k-1} seconds deterministically. That means that the frequency of switching is 2^{-k+1} for the kth bit, which corresponds to the "critical" decay in BF model. Let us see what we have in first few steps of the process:

	$k=1$	$k=2$	$k=3$	$k=4$	$\dots$
$t=0$	0	0	0	0	$\dots$
$t=1$	1	0	0	0	$\dots$
$t=2$	0	1	0	0	$\dots$
$t=3$	1	1	0	0	$\dots$
$t=4$	0	0	1	0	$\dots$
$t=5$	1	0	1	0	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$

The sequence of bits at time $t = 0, 1, 2, \dots$ is exactly the binary representation of t. Therefore, no less than one bit is active at all times. The

same is true for slower decay of frequencies: the next bit always turns on before the previous turns off. However, if we now increase the decay speed, make it $(2+\varepsilon)^{-k+1}$ for some $\varepsilon > 0$, then the first k bits will have time to go through all possible combinations (including all of the bits being idle) before the $k+1$-st bit becomes active.

Theorem 3.2. *When a BF model is recurrent, it is null-recurrent, i.e.* $\mathbf{E}\,\tau_{\mathrm{BF}} = \infty$.

Theorem 3.3. *Consider the recurrent BF model in discrete time with* $p_k \sim C_1 p^k$ *for some fixed constant* $C_1 > 0$ *and* $p \in (0, 1/2)$. *Then*

(i) $\mathbf{E}\,\tau_{\mathrm{BF}}^r < \infty$ *for any positive* $r < 1 - \frac{\log 2}{\log(1/p)}$. *Moreover, for any such* r, *if the Markov chain* (1) *is started from an arbitrary* $\zeta_0 \in \mathcal{X}$ *with the largest active bit* M_0, *then there exists a constant* $C_2 = C_2(C_1, p, r)$ *such that*

$$\mathbf{E}\,[\tau_{\mathrm{BF}}^r | M_0 = m] \le C_2 \left(\frac{1}{2p}\right)^m;$$

(ii) $\mathbf{E}\,\tau_{\mathrm{BF}}^r = \infty$ *for any* $r > 1 - \frac{\log(2-p)}{\log(1/p)}$.

Next we show that the DB models can be recurrent or transient, depending on the decay of p_k's. Define a tail of distribution $\mathcal{P}$ as

$$Q_k = \sum_{j=k+1}^{\infty} p_k.$$

Theorem 3.4. *If the distribution* $\mathcal{P}$ *is such that:*

(i) $\limsup_{k\to\infty} \frac{Q_{k+1}}{Q_k} = p < 1$, *then the DB model is recurrent,*

(ii) $p_k \sim C\exp(-\alpha k^\gamma), k \to \infty$ *for some* $\alpha > 0, \gamma \in (0, 1/2)$, *then the DB model is transient.*

Finally, we prove a central limit theorem for η_t – the number of active bits in the continuous-time version of BF, DB models.

Theorem 3.5. *For both BF and DB models, whenever*

$$\mathbf{E}\,\eta_t \to \infty, \tag{2}$$

then also $\mathbf{var}\,\eta_t \to \infty$ *as* $t \to \infty$ *and*

$$\frac{\eta_t - \mathbf{E}\,\eta_t}{\sqrt{\mathbf{var}\,\eta_t}} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1) \quad \text{as } t \to \infty.$$

In BF model the condition (2) *is always fulfilled, and in DB model a sufficient condition for* (2) *is:*

$$p_k \sim C \exp(-\alpha k^\gamma), k \to \infty, \tag{3}$$

for some constants $C > 0, \alpha > 0, \gamma \in (0, 1)$.

The proofs can be found in paper II.

4 Random exchange models

4.1 Motivation: random subsequent triangulation

The original motivation for studying the random exchange models comes from the triangulation models [21]. In [7], authors consider the barycentric triangulation, where at each step the triangle is divided by its medians into 6 new triangles, then one of the new triangles is chosen uniformly at random, and this procedure is iterated. The object of interest is the limiting shape of the triangles. In [7] authors prove that the triangles weakly converge to the flat shape, which seems to be impractical for applications. In [21] authors prove similar results for subdivision with bisectors and subdivision where the middle point is chosen uniformly inside the triangle at each step.

In this connection, it is interesting to consider the iterations of the following procedure. Start with a triangle, draw its circumcircle. The vertices of the triangle split the circle into three arcs. Sample three new points: one in each arc, uniformly distributed inside of the arc as shown on Figure 3. Connect the new points and erase the old triangle. We are now left with a new triangle, obtained from an old one in a random fashion, similarly to the model in [21]. It is possible to prove that the weak limit of this new procedure is a non-degenerate random triangle with an explicitly known distribution of the angles.

The shape of a triangle is defined by the magnitudes of its three angles. The angles are proportional to the lengths of the opposing arcs of the circumcircle, so we can limit our analysis to the arc lengths' proportions. During the proposed procedure, each of the new arcs is composed of the two uniformly random chunks of the old ones. If τ_1, τ_2, τ_3 are the lengths of the old arcs, and $\tau_1', \tau_2', \tau_3'$ are the lengths of the new ones, then we can write it in the matrix form:

$$(\tau_1', \tau_2', \tau_3') = (\tau_1, \tau_2, \tau_3) \begin{pmatrix} 1-U_1 & U_1 & 0 \\ 0 & 1-U_2 & U_2 \\ U_3 & 0 & 1-U_3 \end{pmatrix} \tag{4}$$

where U_1, U_2, U_3 are independent random variables, uniformly distributed on $(0, 1)$. This procedure of a selection of a new triangle then falls into the

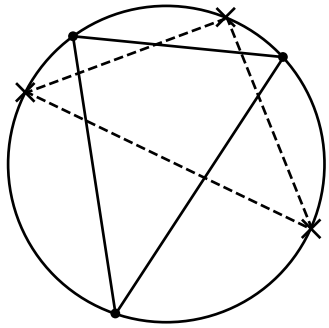

Figure 3: Vertices of the new triangle (dashed) are uniformly selected on the arcs between the vertices of the original one (solid).

framework of random mass exchange models, as described in [17, Section 4]. The results of [17] heavily depend on the properties of the Dirichlet distributions, which we state below.

4.2 Dirichlet distributions

Definition 4.1. A random variable η has a Gamma distribution with parameters $\alpha, \lambda > 0$ (shape, rate) if its density $f(x)$ is given by

$$f(x) = \begin{cases} \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x}, & x > 0. \\ 0, & \text{otherwise} \end{cases}$$

We will sometimes omit the rate parameter λ, defaulting it to 1.

Definition 4.2. A random vector $X = (X_1, \ldots, X_r)$ with support on an r-dimensional, $r \geq 2$, simplex

$$T = \left\{ (x_1, x_2, \ldots, x_r) : \sum_{k=1}^{r} x_k = 1,\ x_k > 0, k = 1, \ldots, r \right\}$$

has Dirichlet distribution with positive real parameters $(\alpha_1, \alpha_2, \ldots, \alpha_r)$ if its density is given by

$$f_X(x) = \frac{\Gamma(\sum_{i=1}^r \alpha_i)}{\prod_{i=1}^r \Gamma(\alpha_i)} \prod_{i=1}^{r} x_i^{\alpha_i - 1},\ x \in T.$$

We write $X \sim \text{Dirichlet}(\alpha_1, \alpha_2, \ldots, \alpha_r)$.

In case $r = 2$, one can directly verify that if a random variable U has distribution $\text{Beta}(\alpha_1, \alpha_2)$, then the random vector $(U, 1-U)$ has distribution $\text{Dirichlet}(\alpha_1, \alpha_2)$. In particular, $\alpha_1 = \alpha_2 = 1$ yields U uniformly distributed on $(0, 1)$, and so the rows of the random exchange matrix in (4) are Dirichlet distributed (with a slight generalization, allowing for $\alpha_k = 0$, with the corresponding vector components equal to 0 almost surely).

It is good to have some intuition for Dirichlet random vectors before formulating the results. The first probabilistic intuition on Dirichlet random vectors comes from seeing them as shapes of vectors of independent Gamma-distributed random variables.

Proposition 4.1 (Probabilistic definition)**.** *If independent random variables $\eta_1, \eta_2, \ldots, \eta_r$ have Γ-distributions with shape parameters $\alpha_1, \alpha_2, \ldots, \alpha_r$ correspondently, and a common rate λ; and $V = \sum_{i=1}^{r} \eta_i$, then*

$$(\eta_1/V, \eta_2/V, \ldots, \eta_r/V) \sim \text{Dirichlet}(\alpha_1, \alpha_2, \ldots, \alpha_r).$$

Another interesting perspective on Dirichlet random vectors comes from treating them as limiting colour proportions in multicolour Polya urn models. In particular, if an r-coloured Polya urn process with a single reinforcement starts with $\alpha_k > 0$ balls of kth colour, $k = 1, 2, \ldots, r$, then the almost sure limit of the proportions of the balls of different colours is a random vector with distribution $\text{Dirichlet}(\alpha_1, \alpha_2, \ldots, \alpha_r)$. An outstanding constructive illustration of this fact and beyond is given in a paper of Blackwell and McQueen [4], already mentioned in Section 2.

We will need the following two simple properties of Dirichlet distributions that easily follow from the two probabilistic interpretations of Dirichlet distributions that we already have.

Proposition 4.2 (Properties of Dirichlet distributions)**.**

1. *(Agglomeration) If $(X_1, \ldots, X_r) \sim \text{Dirichlet}(\alpha_1, \ldots, \alpha_r)$, then*

$$(X_1, \ldots, X_{k-1}, X_k + X_{k+1}, X_{k+2}, \ldots, X_r) \\ \sim \text{Dirichlet}(\alpha_1, \alpha_{k-1}, \alpha_k + \alpha_{k+1}, \alpha_{k+2}, \ldots, \alpha_r).$$

2. *(Decimation) If*

$$X = (X_1, \ldots, X_r) \sim \text{Dirichlet}(\alpha_1, \ldots, \alpha_r)$$

and

$$Y = (Y_1, \ldots, Y_r) \sim \text{Dirichlet}(\delta_1, \ldots, \delta_s)$$

where $\sum_{j=1}^{s}\delta_j = \alpha_k$, *$X$ and Y independent, then*

$$(X_1,\dots,X_{k-1},\quad X_kY_1, X_kY_2,\dots,X_kY_s,\quad X_{k+1},X_{k+2},\dots,X_r)$$

$$\sim \text{Dirichlet}(\alpha_1,\dots,\alpha_{k-1},\quad \alpha_k\delta_1,\alpha_k\delta_2,\dots,\alpha_k\delta_s,\quad \alpha_{k+1},\dots,\alpha_r).$$

The last and the most important property that we need in our context of random exchange models, is tying together the Dirichlet and Gamma distributions. We use the shape versus size independence characterisation of Gamma random vectors, obtained in [18] as an extension of the main result of [16].

Proposition 4.3 (Shape versus size independence)**.** *Given the positive, non-degenerate independent random variables $Y_1,\dots,Y_r$, if $V=\sum_{k=1}^{r}Y_k$ is the size random variable, then the shape vector $(Y_1/V,\dots,Y_r/V)$ is independent of V if and only if for every k, Y_k has a Gamma distribution with parameters $\alpha_k>0,\lambda>0$. In that case, the shape vector has Dirichlet distribution with parameters $(\alpha_1,\dots,\alpha_r)$*

A straightforward corollary from this property allows us to build vectors of independent Gamma-distributed random variables from Dirichlet random vectors.

Proposition 4.4. *Assume $V\sim\Gamma(\alpha,\lambda)$ is independent of $(X_1,\dots,X_r)\sim$ Dirichlet$(\alpha_1,\dots,\alpha_r)$, where $\sum_{k=1}^{r}\alpha_k=\alpha$. Then the random variables $Y_k=X_kV, k=1,\dots,r$, are jointly independent, and $Y_k\sim\Gamma(\alpha_k,\lambda), k=1,\dots,r$.*

4.3 Random exchange models

The procedure of selection of a new triangle defined in (4) can be thought of as a random mass exchange. Let us think of the three vertices of the triangle as three agents, each holding a certain amount of mass, which in this case corresponds to the value of the angle near the respective vertex. At each step, each agent passes a uniform proportion of mass at hand to their neighbour in the counter-clockwise direction, or to the left, while at the same time receiving a portion of mass from their neighbour in the clockwise direction, or to the right. If we denote by $\Pi(n)$ the random exchange matrix holding in each row the vector of sharing proportions, then iterating the procedure amounts to a right multiplication: if τ^n denotes the vector of angles after n iterations, we can write

$$\tau^n=\tau^0\Pi(1)\Pi(2)\dots\Pi(n)=:\tau^0\Pi(1:n) \tag{5}$$

It is known (see [5], [17]) that under a certain set of conditions, the almost sure limit as $n\to\infty$ of the left product $\Pi(n:1)$ of random

matrices with independent Dirichlet rows exists and is a matrix with a.s. identical Dirichlet rows. Therefore, since $\Pi(1:n) \stackrel{\mathcal{D}}{=} \Pi(n:1)$ for $n = 1, 2, \ldots$, we conclude that the distributional limit of (5) exists and is the same for all initial triangle shapes, defined by the vector of angles τ^0. The limiting distribution of the angles vector, divided by π, the sum of angles of the triangle is Dirichlet$(2, 2, 2)$, as it is left invariant under the dynamics (4), as easily follows from the properties of Dirichlet distributions stated in Proposition 4.2.

Random exchange in high dimensions. Coming back to the arc-length interpretation, it is interesting to allow for more than 3 arcs at a time on the circle. Assume that we start with d equidistant points on the circle, dividing it into d equal arcs. At each step n, let each point move to a new location, which is uniformly chosen on the arc to the right from it. Iterating this procedure, we will then arrive to the limiting distribution of arcs' lengths that will be Dirichlet$(2, \ldots, 2)$ – following the same logic.

4.4 Main results

We consider the analogue of the above dynamics when the circle is replaced by the real line with an infinite amount of points on it. We define a neighbour-dependent random shift procedure for a stationary renewal process on the real line. Given a stationary renewal process $\mathbb{T} = \{\ldots < T_{-1} < T_0 < 0 < T_1 < T_2 < \ldots\}$, with jump distribution F and the correspondent sequence of increments $\mathcal{T} = \{\tau_k\}_{k\in\mathbb{Z}}$, $\tau_k = T_k - T_{k-1}, k \in \mathbb{Z}$, define a stochastic shift operator $\Psi_G(\mathbb{T})$ as follows. Let $\{b_k, k \in \mathbb{Z}\}$ be i.i.d. random variables with distribution G, and $\Psi_G(\mathbb{T}) = \mathbb{T}' = \{\ldots < T'_{-1} < T'_0 < T'_1 < T'_2 < \ldots\}$ be the result of the random shift operation where each point is moved to a random G-distributed location inside of the interval to the right from it, as shown on Figure 4:

$$T'_k = T_k + b_k(T_{k+1} - T_k), k \in \mathbb{Z}$$

The shift operator Ψ_G induces a random exchange operator Φ_G on the sequence of interval lengths:

$$\Phi_G(\mathcal{T}) = \mathcal{T}' = \{\tau'_k\}_{k\in\mathbb{Z}}, \text{ where } \tau'_k = (1 - b_k)\tau_k + b_{k+1}\tau_{k+1}, k \in \mathbb{Z}. \quad (6)$$

So kth interval passes a b_k proportion of its length to its neighbour to the left, while at the same time receiving a b_{k+1} proportion of the length of its neighbour to the right.

It is clear that if $\mathbb{T}$ is a stationary point process, then the result of the operation $\Psi_G(\mathbb{T})$ is also a stationary point process. However, the sequence $\mathcal{T}$ is not i.i.d. in general, due to the size bias of the length of the interval τ_1 containing the origin. Besides, the interval (T'_0, T'_1) does

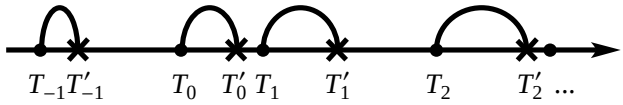

Figure 4: Neighbour-dependent shift operator Ψ_G: every point of $\mathbb{T}$ (solid dots) is moved to a G-distributed random location inside of the interval to the right from it (crosses).

not necessarily contain the origin, so we cannot think of $\mathbb{T}$ as a stationary renewal process without changing the indices in a suitable manner. In order to deal with these two issues, it is convenient to think in terms of Palm distributions of the correspondent point processes. Under Palm distribution of $\mathbb{T}$ there is almost surely a point $T_0 = 0$ of the process in the origin, and the sequence of increments $\mathcal{T} = \{\tau_k\}$ is i.i.d. with distribution F. Under Palm distribution of $\mathbb{T}'$ there is also a point in the origin, and the sequence of increments is equal in distribution to $\{\tau'_k\}_{k\in\mathbb{Z}} = \mathcal{T}' = \Phi_G(\mathcal{T})$ defined as in (6), but with τ_k now independent, identically distributed as F. Our first result is the characterisation of the class of distributional fixed points of the exchange operator Φ_G. Denote by δ_x the degenerate distribution concentrated in x. Obviously, when G is δ_0 or δ_1, the corresponding Φ_G preserves the distribution of any stationary sequence, in particular $\mathcal{T}$, so we exclude these trivial cases from our consideration.

Theorem 4.1. *Let $\mathcal{T}$ be an i.i.d. sequence of positive integrable random variables with the distribution F corresponding to the sequence of increments under the Palm distribution of the stationary renewal process $\mathbb{T}$, and Φ_G be the random exchange operator defined as in* (6)*. Then $\Phi_G(\mathcal{T}) \stackrel{\mathcal{D}}{=} \mathcal{T}$ (and hence $\Psi_G(\mathbb{T}) \stackrel{\mathcal{D}}{=} \mathbb{T}$) if and only if one of the alternatives is true:*

(i) $F = \Gamma(\alpha, \gamma)$ and $G = \mathrm{Beta}(r\alpha, (1-r)\alpha)$ for some constants $\alpha > 0$, $\gamma > 0$ and $r \in (0, 1)$,

(ii) $F = \delta_s$ for some $s \in (0, \infty)$ and $G = \delta_b$ for some $b \in (0, 1)$.

Corollary 4.1. *Since a homogeneous Poisson process with rate γ has exponential $\Gamma(1, \gamma)$ distributed interpoint distances, the Beta division point distribution $G = B(r, 1-r)$ for some $0 < r < 1$ is the only non-degenerate distribution for which the random shift Ψ_G preserves the Poisson process.*

Corollary 4.1 seems counter-intuitive: each point is moved to a location that heavily depends on its neighbour's position, so the interval

lengths of the shifted process $\mathbb{T}'$ seem to be 2-dependent, breaking the definition of a Poisson process. However, the unique joint properties of Gamma and Beta distributions imply that the lengths of the parts $b_k\tau_k$ and $(1-b_k)\tau_k$ in which the point of a new process T'_k divides the kth interval are independent. Paradoxically, we can cut a Gamma-distributed interval in Beta-distributed proportions in two independent pieces!

Corollary 4.2. *Every second point in a homogeneous Poisson process form a renewal process with $\Gamma(2,\gamma)$-distributed interpoint distances. Thus a uniform division distribution which is also $G = B(1,1)$ preserves it. This also follows from a known elementary fact that if X, Y are independent Exponentially-distributed random variables and U is a uniform variable independent of them, then $U(X+Y)$ and $(1-U)(X+Y)$ are independent Exponentially-distributed random variables.*

Corollary 4.3. *The random adjustment procedure, moving every second point to a uniform location between its two neighbours, preserves the Poisson process.*

Graphical intuition. The following embedding of the Gamma renewal process on the positive half of the real line into a Levy process with Gamma increments provides an illustration of the "if" part of Theorem 4.1. Let positive α, γ and $0 < r < 1$ be fixed. Consider a Gamma process $Y(t), t \in [0,\infty)$ which is a strictly increasing Lévy process with Gamma-distributed increments, so that

- $Y(0) = 0$, and
- for any n, $0 \le t_0 < t_1 < t_2 < \ldots < t_n$, the random variables $Y(t_1)-Y(t_0), \ldots, Y(t_n)-Y(t_{n-1})$ are independent, Gamma-distributed with a common rate parameter γ, and shape parameters $(t_1-t_0)\alpha, \ldots, (t_n-t_{n-1})\alpha$, respectively.

See [9] for a constructive definition of such $Y(t)$. Then, for $k = 0, 1, 2, \ldots$, put $T_k = Y(k)$ and $T'_k = Y(k+r) - Y(r)$, so that the increments are defined by

$$\begin{aligned}\tau_k &= Y(k) - Y(k-1),\\ \tau'_k &= Y(k+r) - Y(k-1+r),\end{aligned}$$

see Figure 5. By construction, the sequence $\mathcal{T} = \{Y(1)-Y(0), Y(2)-Y(1), Y(3)-Y(2), \ldots\}$ is is a sequence of independent random variables with distribution $\Gamma(\alpha,\gamma)$. So is the sequence $\mathcal{T}' = \{Y(1+r)-Y(r), Y(2+r)-Y(1+r), Y(3+r)-Y(2+r), \ldots\}$. Show that the latter is obtained from the former as a result of random shift procedure with a suitable i.i.d. sequence $\{b_k\}$ with Beta$(r\alpha, (1-r)\alpha)$ distribution.

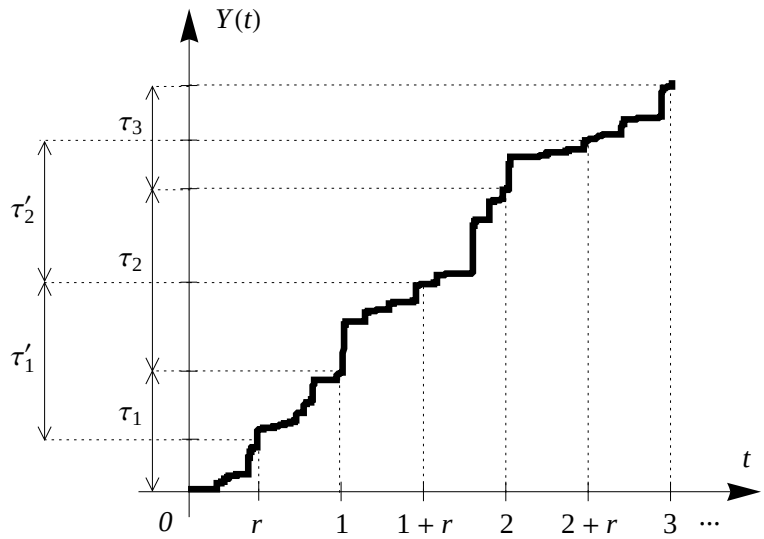

Figure 5: Gamma process $Y(t)$ and the coupling of $\mathcal{T}$ and $\mathcal{T}'$.

Since

$$\begin{aligned} Y(k+r) - Y(k-1+r) =& Y(k) - Y(k-1+r) + Y(k+r) - Y(k) \\ =& \frac{Y(k) - Y(k-1+r)}{Y(k) - Y(k-1)}(Y(k) - Y(k-1)) \\ &+ \frac{Y(k+r) - Y(k)}{Y(k+1) - Y(k)}(Y(k+1) - Y(k)), \end{aligned}$$

define

$$b_k = \frac{Y(k-1+r) - Y(k-1)}{Y(k) - Y(k-1)}, \quad k = 1, 2, \ldots$$

Notice that the denominator

$$Y(k) - Y(k-1) = (Y(k) - Y(k-1+r)) + (Y(k-1+r) - Y(k-1))$$

is the sum of two independent Gamma-distributed random variables, so b_k's are independent for different k and $\mathrm{Beta}(r\alpha, (1-r)\alpha)$-distributed. Moreover, by the shape vs. size independence property of the Gamma distribution, the sequence $\{b_k\}_{k\geq 1}$ is independent of $\{\tau_k\}_{k\geq 1}$.

Convergence. The next question of interest is the convergence of iterations of Φ_G. We now move on to the random mass exchange setting.

Consider a countable collection of agents labelled by integers $\mathbb{Z}$. At time $n = 0$, kth agent holds an initial amount τ_k^0 of mass, where $\{\tau_i^0\}_{i\in\mathbb{Z}}$ is a sequence satisfying two conditions:

(A1) *The initial masses τ_i^0 are non-negative i.i.d. random variables for different $i \in \mathbb{Z}$, with finite first moment $\mathbf{E}\,\tau_i^0 < \infty$*

(A2) *The second moments of initial masses are finite:* $\mathbf{var}\,\tau_i^0 = \sigma^2 < \infty$

Define the random exchange process iteratively. At the beginning of step $n = 1, 2, \dots$ the ith agent has mass τ_i^{n-1}. Then ith agent samples a new vector of proportions $(\pi_{i,i+j}(n))_j \in \mathbb{Z}$ and distributes all of its mass between itself and other agents accordingly, so that $(i+j)$th agent receives an amount $\tau_i^n \pi_{i,i+j}(n)$, so the new row-vector of masses τ^n can be obtained from an old one by right-multiplication by a random exchange matrix $\Pi(n)$:

$$\tau^n = \tau^{n-1}\Pi(n) = \tau^0\Pi(1:n),$$

where $\Pi(n) = (\pi_{ij}(n))_{i,j\in\mathbb{Z}}$ contains the exchange proportions, and $\Pi(1:n) = \Pi(1)\dots\Pi(n)$ is a shorthand for right multiplication.

We assume that the random exchange model is time-homogeneous, and translation-invariant on $\mathbb{Z}$:

(B1) *There exists a random probability distribution $(\pi_j)_{j\in\mathbb{Z}}$ on $\mathbb{Z}$ such that the vectors $(\pi_{i,i+j}(n))_{j\in\mathbb{Z}}$ for different $i \in \mathbb{Z}$, and $n = 1, 2, \dots$ are i.i.d. copies of $(\pi_j)_{j\in\mathbb{Z}}$*

Theorem 4.2. *Let the sharing proportion matrices $\Pi(n), n = 1, 2, \dots$ satisfy the condition (B1). Then there exists a unique (up to a distributional copy and scaling by a constant factor) fixed point, i.e. a vector of random variables (not necessarily independent) $\tau^\infty = (\tau_i^\infty)_{i\in\mathbb{Z}}$, such that:*

(i) $\tau^\infty\Pi(n) \stackrel{\mathcal{D}}{=} \tau^\infty$,

(ii) for any vector of initial masses τ^0 satisfying (A1-A2), $\tau^0\Pi(1:n) \Rightarrow \tau^\infty, n \to \infty$.

Our last result characterises the random exchange processes for which the limit configurations have independent coordinates, as it is in the case of a random exchange process corresponding to neighbor-dependent shift Ψ_G of a renewal process defined in the beginning; in a special case where the proportion vector π is exchangeable.

Theorem 4.3. *Assume $\tau = (\tau_j)_{j\in\mathbb{Z}}$, Π are independent of each other, satisfying (A1-A2), (B2). Assume additionally, that only finite amount of vector π's coordinates in (B2) are non-zero:*

$$|\{i \in \mathbb{Z} : \mathbf{P}(\pi_i > 0) > 0\}| < \infty,$$

and that the distribution of those coordinates is exchangeable. Put $\tau' = \tau\Pi$. Then

$$\tau' \stackrel{\mathcal{D}}{=} \tau$$

if and only if the components of τ are Gamma-distributed: $\tau_j \sim \Gamma(a, \gamma)$, and the rows of Π are Dirichlet-distributed: $(\pi_i)_{i\in K} \sim \mathrm{Dirichlet}((ar)_{i\in K})$, where $r = 1/|K|$ and $a, \gamma > 0$.

5 References

[1] K. Athreya and S. Karlin. Embedding of urn schemes into continuous time Markov branching process and related limit theorem. *Ann. Math. Statist.*, 39:1801–1817, 1968.

[2] J. Balogh and R. Pemantle. The Klee-Minty random edge chain moves with linear speed. *Random Struct. Algor.*, 30(4):371–390, 2007.

[3] P. Berti, I. Crimaldi, L. Pratelli, and P. Rigo. A central limit theorem and its applications to multicolor randomly reinforced urns. *J. Appl. Probab.*, 48(2):527–546, 2011.

[4] D. Blackwell and J. MacQueen. Ferguson distributions via Polya urn schemes. *Ann. Stat.*, 1:353–355, 1973.

[5] J.-F. Chamayou and G. Letac. A transient random walk on stochastic matrices with Dirichlet distributions. *Ann. Probab.*, 22(1):424–430, January 1994.

[6] F. Chung, S. Handjani, and D. Jungreis. Generalizations of Polya's urn problem. *Ann. Combin.*, 7:141–153, 2003.

[7] P. Diaconis and L. Miclo. On barycentric subdivision. *Comb. Probab. Comput.*, 20(2):213–237, March 2011.

[8] L. Dubins and D. Freedman. Random distribution functions. *B. Am. Math. Soc.*, 69:548–551, 1963.

[9] T. Ferguson and M. Klass. A representation of independent increment processes without Gaussian components. *Ann. Math. Statist.*, 43:1634–1643, 1972.

[10] P. Flajolet, P. Dumas, and V. Puyhaubert. Some exactly solvable models of urn process theory. In *proc. of 4th Colloquium on Mathematics and Computer Science (DMTCS)*, pages 59–118. Nancy, France, 2006.

[11] R. Gouet. Martingale functional central limit theorems for a generalized Polya urn. *Ann. Prob.*, 21(3):1624–1639, 1993.

[12] O. Häggström, Y. Peres, and J. Steif. Dynamical percolation. *Annales de L'I.H.P., section B*, 33(4):497–528, 1997.

[13] K. Inoue and S. Aki. A generalized Pólya urn model and related multivariate distributions. *Ann. Inst. Statist. Math.*, 57(1):49–59, 2005.

[14] S. Kakutani. A problem on equidistribution on the unit interval [0,1]. *Measure Theory*, 541:369–375, 1975.

[15] S. Kotz, H. Mahmoud, and P. Robert. On generalized Pólya urn models. *Stat. Prob. Letters*, 49:163–173, 2000.

[16] Eugene Lukacs. A characterization of the gamma distribution. *Ann. Math. Stat.*, 26(2):319–324, 06 1955.

[17] Shaun McKinlay. A characterisation of transient random walks on stochastic matrices with dirichlet distributed limits. *arXiv preprint arXiv:1212.4975*, 2012.

[18] James E Mosimann. Size allometry: size and shape variables with characterizations of the lognormal and generalized gamma distributions. *J. Am. Stat. Ass.*, 65(330):930–945, 1970.

[19] R. Pemantle. A survey of random processes with reinforcement. *Prob. Surveys*, 4:1–79, 2007.

[20] J. Peyriére. A singular measure generated by splitting [0,1]. *Z. Wahrsch. verw. Gebiete*, 47:289–297, 1979.

[21] Stanislav Volkov. Random geometric subdivisions. *Random Struct. Algor.*, 43(1):115–130, 2013.

[22] S. Zuyev. Stopping sets: Gamma-type results and hitting properties. *Adv. Appl. Prob.*, 31(2):355–366, 1999.