

Probabilistic Non-asymptotic Analysis of Distributed Algorithms

Nicolas Champagnat*, René Schott†and Denis Vilemonais‡

Keywords: Distributed algorithms, Deadlock, Quasi-stationary distribution.

Abstract

We present a new probabilistic analysis of distributed algorithms. Our approach relies on the theory of quasi-stationary distributions (QSD) recently developed by Champagnat and Vilemonais [4, 5, 6]. We give properties on the deadlock time and the distribution of the model before deadlock, both for discrete and diffusion models. Our results are non-asymptotic since they apply to any finite values of the involved parameters (time, numbers of resources, number of processors, etc.) and reflect the real behavior of these algorithms, with potential applications to deadlock prevention, which are very important for real world applications in computer science.

1 Introduction

Today's distributed systems involve a huge (but finite) number of processors sharing common resources (i.e. are massively parallel). These systems are inherently fragile. For example, if a processor is running out of memory it can stop the whole system and deadlock may appear. Usually, analysis of distributed systems leads to asymptotic results (where the parameters of interest tend to infinity) and characterization of limit laws (of large numbers,

*Inria and Institut Elie Cartan, Université de Lorraine, Nancy, France

†Institut Elie Cartan and LORIA, Université de Lorraine, Nancy, France

‡Institut Elie Cartan, Université de Lorraine, Nancy, France

central limit theorems, etc.). But infinity does not exist in computer science! Data structures have large but finite dimension, even the most efficient computer is not able to realize an infinite number of operations. Asymptotic results are, therefore, not very useful. Non-asymptotic results are true for any value of the parameters of interest (time, memory area, etc.) and are, therefore, closer to real world applications in computer science. The purpose of this paper is to present a non-asymptotic analysis of distributed algorithms which works for any finite number of processors and any finite number of resources (a similar study is possible for other types of algorithms).

Our approach relies on the theory of quasi-stationary distributions (QSD) recently developed by Champagnat, Villemonais et al. [3, 4, 5, 6, 21].

The organization of this paper is as follows: simple distributed algorithms are presented in Section 2. Section 3 contains generalities on quasi-stationary distributions in finite state spaces. Our main results are stated and proved in Section 4. Simulations are presented in Section 5. Section 6 contains concluding remarks and some further research aspects.

2 Examples

We describe two models of distributed systems with possible deadlocks, which will be studied numerically in Section 5.

2.1 Colliding stacks

The presentation of this example is based on [14]. For pedagogical reasons, we consider only two stacks but, of course, real storage allocation algorithms involve a huge number of stacks.

Assume that two stacks are to be maintained inside a shared (contiguous) memory area of a fixed size m . A trivial algorithm will let them grow from both ends of that memory area until their cumulative sizes fill the initially allocated storage (m cells), and the algorithm stops having exhausted the available memory. That shared storage allocation algorithm is to be compared to another option, namely allocating separate zones of size $m/2$ to each of the two stacks. This separate storage allocation method will then halt as soon as any one of the two stacks reaches sizes $m/2$. Several measures may be introduced to compare these two schemes. One of them is the number of operations that can be treated by the algorithms under some appropri-

ate probabilistic model. Another interesting measure of the efficiency of the shared allocation that was proposed by Knuth [16], is the size of the largest stack when both stacks meet and the algorithm runs out of storage. Flajolet [14] completely analyzed (combinatorially) this problem and thus solved a question posed by Knuth [16] (Vol. 1, Exercice 2.2.2.13). Partial results have been obtained earlier by Yao [25], but it appears that covering all cases of the original problem cannot be achieved by an extension of Yao's methods. As has been noticed since the problem was initially posed by Knuth [16], the natural formulation is in terms of random walks. Here the random walk takes place in a triangle in a 2-dimensional lattice space: a state is the couple formed with the size of both stacks. The random walk has two reflecting barriers along the axes (a deletion takes no effect on an empty stack) and one absorbing barrier parallel to the second diagonal (the algorithm stops when the combined sizes of the stacks exhaust the available storage). Probabilistic analyses of the colliding stacks problem have been done (under various assumptions) by Louchard and Schott [18], Louchard, Schott, Tolley and Zimmermann [17], Maier [19], Guilloin-Plantard and Schott [15], Comets, Delarue and Schott [9, 10].

2.2 Banker algorithm

For simplicity, we restrict the presentation (as for the colliding stacks) to dimension 2.

Consider two customers P_1 and P_2 sharing a fixed quantity of a given resource R (money, say). There are fixed upper bounds m_1, m_2 on how much of the resource each of the customers will need at any time. The banker decides to affect to the customer P_i ($i = 1, 2$) the required units only if the remaining units are sufficient in order to fulfill the requirements of P_j ($j = 1, 2; j \neq i$). The situation is modeled by a random walk in a rectangle with a broken corner (i.e. $(x_1, x_2) : 0 \leq x_1 \leq m_1, 0 \leq x_2 \leq m_2, x_1 + x_2 \leq m$, where the last constraint generates the broken corner). The random walk is reflected on the sides parallel to the axes and is absorbed on the sloping side.

Probabilistic analyses of this algorithm have been presented in [18, 17, 15] for two customers and in [9, 10] for $d \in N$ customers.

Maier and Schott [20] proved partial results for d customers and $r \in N$ resources.

2.3 Description of the Model in higher dimension

We consider the interaction of q processes P_1, P_2, \dots, P_q , each with its own resource needs. We allow the processes to access to r different, non-substitutable resources (i.e. types of memory) R_1, R_2, \dots, R_r . We model resource limitations, and define resource exhaustion as follows. At any time s , process P_i is assumed to have allocated some quantity $y_i^j(s)$ of resource R_j . Process P_i is assumed to have some maximum need m_{ij} of resource R_j so that

$$0 \leq y_i^j \leq m_{ij}, \quad \forall s \geq 0. \quad (1)$$

The constant m_{ij} may be infinite; if finite, it is a hard limit which the process P_i never attempts to exceed. The resources R_j are limited:

$$\sum_{i=1}^q y_i^j(s) < m_j, \quad (2)$$

so that $m_j - 1$ is the total amount of resource R_j available for allocation. Remember that resource exhaustion occurs when some process P_i issues an unfulfillable request for a quantity of some resource R_j . Here "unfulfillable" means that fulfilling the request would violate one of the inequalities (2).

The state space of the memory allocation system is a convex polytope with faces defined by hyperplanes H_1, \dots, H_k which are either reflective or absorbing.

- For example, we can consider a model with d processors and a single resource with limit m , so that the state space of the random walk is

$$E = \{x \in [0, a_1] \times \dots \times [0, a_d] \text{ s.t. } \langle x, v \rangle < m\},$$

where v is a vector of $\mathbb{R}_+^d \setminus \{0\}$ with nonnegative coordinates and with absorbing state

$$\partial = \{x \in [0, a_1] \times \dots \times [0, a_d] \text{ s.t. } \langle x, v \rangle = m\}.$$

- General case. For d processors and r resources,

$$E = \{x \in [0, a_1] \times \dots \times [0, a_d] \text{ s.t. } \langle x, v_j \rangle < m_j, \forall j = 1, \dots, r\}, \quad (3)$$

where v_j are vectors of $(\mathbb{R}_+)^d \setminus \{0\}$ with nonnegative coordinates and m_j is the maximum amount of resource j , and with absorbing state

$$\partial = \{x \in [0, a_1] \times \dots \times [0, a_d] \text{ s.t. } \langle x, v_j \rangle = m_j, \forall j = 1, \dots, r\}. \quad (4)$$

3 Quasi-stationary distributions

The goal of this section is to give a short survey on the main results on quasi-stationary distributions for absorbed Markov processes and their implications on deadlock prevention and analysis. To keep things simple, we focus here on the case of continuous-time processes taking values in a finite state space, like in Section 2.2.

We consider a Markov process $(X_t, t \geq 0)$ taking values in a finite state space $E \cup \partial$, where ∂ is absorbing, meaning that $X_t \in \partial$ a.s. for all $t \geq \tau_\partial := \inf\{t \geq 0 : X_t \in \partial\}$. We assume that $\tau_\partial < \infty$ a.s., i.e. ∂ is accessible from any state in E . In the context of the colliding stacks and the banker models, τ_∂ is the deadlock time. For all $x \neq y \in E \cup \partial$, we denote by $q_{x,y} \geq 0$ the transition rate from x to y and we set as usual

$$q_{x,x} := -q_x := \sum_{y \neq x} q_{x,y}.$$

We assume in all this section that the matrix $Q := (q_{x,y})_{x,y \in E}$ is irreducible, so that Perron-Frobenius theorem applies to the exponential of the matrix Q : when $t \rightarrow +\infty$,

$$(e^{tQ})_{x,y} = e^{-\lambda_0 t} v_x u_y + O(e^{-\lambda_1 t}), \quad (5)$$

where $-\lambda_0$ is the spectral radius of the matrix Q , u and v are the normalized, positive left and right eigenvector of Q for the eigenvalue $-\lambda_0$, i.e. $uQ = -\lambda_0 u$, $Qv = -\lambda_0 v$, $\sum_x u_x = 1$ and $\sum_x u_x v_x = 1$, and

$$-\lambda_1 := \sup_{\lambda \in \text{Sp}(Q), \lambda \neq \lambda_0} \text{Re}(\lambda),$$

where $\text{Sp}(Q)$ is the spectrum of the matrix Q in \mathbb{C} and $\text{Re}(z)$ is the real part of $z \in \mathbb{C}$. Note that, since the matrix Q is irreducible and sub-conservative in the sense that $Q\mathbf{1} \leq 0$ with at least one negative coordinate, where $\mathbf{1}$ is the vector of \mathbb{R}^E with all coordinates equal to 1, Perron-Frobenius theory entails that $-\lambda_1 < -\lambda_0 < 0$.

The next result was first proved in [11] and follows easily from the formulas

$$\mathbb{P}_x(t < \tau_\partial) = \mathbb{P}_x(X_t \in E) = \sum_{y \in E} (e^{tQ})_{x,y}$$

and

$$\mathbb{P}_x(X_t = y \mid t < \tau_\partial) = \frac{\mathbb{P}_x(X_t = y)}{\mathbb{P}_x(X_t \in E)} = \frac{(e^{tQ})_{x,y}}{\sum_{z \in E} (e^{tQ})_{x,z}}.$$

Proposition 3.1. *There exists a constant C such that, for all $x, y \in E$ and all $t \geq 0$,*

$$|\mathbb{P}_x(X_t = y \mid t < \tau_\partial) - u_y| \leq Ce^{-(\lambda_1 - \lambda_0)t} \quad (6)$$

and the probability measure $u = (u_x, x \in E)$ on E is a quasi-stationary distribution, in the sense that

$$\mathbb{P}_u(X_t = y \mid t < \tau_\partial) = u_y, \quad \forall y \in E, t \geq 0, \quad (7)$$

where $\mathbb{P}_u = \sum_{x \in E} u_x \mathbb{P}_x$. In addition, for all $x \in E$ and $t \geq 0$,

$$|e^{\lambda_0 t} \mathbb{P}_x(t < \tau_\partial) - v_x| \leq Ce^{-(\lambda_1 - \lambda_0)t}. \quad (8)$$

Quasi-stationary distributions like u satisfy general properties as explained for example in [21], summarized in the next proposition.

Proposition 3.2. *When the initial population is distributed according to the quasi-stationary distribution u , the absorption time τ_∂ is exponentially distributed with parameter λ_0 , i.e.*

$$\mathbb{P}_u(t < \tau_\partial) = e^{-\lambda_0 t}, \quad \forall t \geq 0,$$

τ_∂ is independent of $(X_{\tau_\partial-}, X_{\tau_\partial})$, where $X_{\tau_\partial-}$ is the position just before exit and X_{τ_∂} is the exit position. In addition, the joint law of $(X_{\tau_\partial-}, X_{\tau_\partial})$ under \mathbb{P}_u is given by

$$\mathbb{P}_u(X_{\tau_\partial-} = x, X_{\tau_\partial} = y) = \frac{u_x q_{x,y}}{\lambda_0}, \quad \forall x \in E, y \in \partial.$$

Note that, as a consequence, the exit position is distributed under \mathbb{P}_u as

$$\mathbb{P}_u(X_{\tau_\partial} = y) = \frac{\sum_{x \in E} u_x q_{x,y}}{\lambda_0}, \quad \forall y \in \partial.$$

This indeed defines a probability distribution on ∂ since

$$\sum_{y \in \partial} \sum_{x \in E} u_x q_{x,y} = \sum_{x \in E} u_x \left(q_{x,x} - \sum_{y \in E, y \neq x} q_{x,y} \right) = \sum_{x \in E} u_x q_{x,x} + \sum_{y \in E} u_y (\lambda_0 - q_{y,y}) = \lambda_0.$$

In addition, the position just before exit is distributed under \mathbb{P}_u as

$$\mathbb{P}_u(X_{\tau_\partial-} = x) = \frac{\sum_{y \in \partial} u_x q_{x,y}}{\lambda_0}, \quad \forall x \in E.$$

This is the quasi-stationary distribution biased by the exit rate $\sum_{y \in \partial} q_{x,y}$ of the process.

Proof. The first property follows from Markov's property and the definition of a quasi-stationary distribution (7)

$$\begin{aligned}\mathbb{P}_u(t + s < \tau_\partial) &= \mathbb{E}_u[\mathbb{1}_{t < \tau_\partial} \mathbb{P}_{X_t}(s < \tau_\partial)] \\ &= \mathbb{E}_u[\mathbb{P}_{X_t}(s < \tau_\partial) \mid t < \tau_\partial] \mathbb{P}_u(t < \tau_\partial) \\ &= \mathbb{P}_u(s < \tau_\partial) \mathbb{P}_u(t < \tau_\partial).\end{aligned}$$

This is the “loss of memory” property characterizing exponential random variables, hence τ_∂ is exponentially distributed.

Since $\mathbb{P}_u(t < \tau_\partial) = \sum_{x \in E} u_x \sum_{y \in E} (e^{tQ})_{x,y}$, it follows from (5) that the parameter of the exponential distribution is λ_0 .

The independence between τ_∂ and X_{τ_∂} follows from a similar computation: for all bounded measurable function f on E ,

$$\begin{aligned}\mathbb{E}_u(f(X_{\tau_\partial}) \mathbb{1}_{t < \tau_\partial}) &= \mathbb{E}_u[\mathbb{E}_{X_t}(f(X_{\tau_\partial})) \mid t < \tau_\partial] \mathbb{P}_u(t < \tau_\partial) \\ &= \mathbb{E}_u(f(X_{\tau_\partial})) \mathbb{P}_u(t < \tau_\partial).\end{aligned}$$

The independence between τ_∂ and $X_{\tau_\partial-}$ can be proved exactly the same way.

Finally, due to the above independence, we have for all $t \geq 0$ and all $x \in E, y \in \partial$,

$$\begin{aligned}\mathbb{P}_u(X_{\tau_\partial-} = x, X_{\tau_\partial} = y) &= \mathbb{P}_u(X_{\tau_\partial-} = x, X_{\tau_\partial} = y \mid \tau_\partial \leq t) \\ &= \frac{\mathbb{P}_u(X_{\tau_\partial-} = x, X_{\tau_\partial} = y, \tau_\partial \leq t)}{\mathbb{P}_u(\tau_\partial \leq t)} \\ &= \frac{\mathbb{P}_u(X_0 = x, X_t = y, J_1 \leq t, J_2 > t) + O(\mathbb{P}_u(J_2 \leq t))}{1 - e^{-\lambda_0 t}},\end{aligned}$$

where $(J_i)_{i \geq 1}$ is the sequence of jump times of the process $(X_t, t \geq 0)$. Standard computations for discrete Markov processes entail

$$\mathbb{P}_u(X_{\tau_\partial-} = x, X_{\tau_\partial} = y) = \frac{tu_x q_{x,y} + O(t^2)}{\lambda_0 t + O(t^2)} \xrightarrow{t \rightarrow 0} \frac{u_x q_{x,y}}{\lambda_0}.$$

This concludes the proof of the Proposition. \square

4 Main results

4.1 Distributions of exit time and exit position

Propositions 3.1. and 3.2. allow us to give estimates on the distributions of the exit time τ_∂ and the exit position $f(X_{\tau_\partial})$ of the process depending on λ_0

and λ_1 . First, (8) entails that

$$\mathbb{P}_x(t < \tau_\partial) \sim v_x e^{-\lambda_0 t} \quad \text{when } t \rightarrow +\infty,$$

hence the distribution of the exit time τ_∂ has an exponential tail. We can also deduce from (8) estimates on the expectation of functions of τ_∂ : for example,

$$\left| \mathbb{E}_x(\tau_\partial) - \frac{v_x}{\lambda_0} \right| \leq \frac{C}{\lambda_1}.$$

Note that this estimate is accurate provided that $\lambda_1 \ll \lambda_0$. This is the regime where the state of the process can be approximated by the quasi-stationary distribution for intermediate times, as explained in the next proposition.

Proposition 4.1. *For all $x \in E$ and $t \geq 0$,*

$$e^{\lambda_0 t} + C e^{-(\lambda_1 - \lambda_0)t} - \frac{e^{\lambda_0 t} - 1 + C e^{-(\lambda_1 - \lambda_0)t}}{u_x} \leq v_x \leq e^{\lambda_0 t} + C e^{-(\lambda_1 - \lambda_0)t}. \quad (9)$$

In addition, for all $x, y \in E$ and $t \geq 0$,

$$|\mathbb{P}_x(X_t = y) - u_y| \leq \frac{u_y}{u_x} (1 - e^{-\lambda_0 t} + C e^{-\lambda_1 t}) + C e^{-\lambda_1 t} (v_x + 2u_y + C e^{-(\lambda_1 - \lambda_0)t}). \quad (10)$$

In the case where $\lambda_0 \ll \lambda_1$, for $1/\lambda_1 \ll t \ll 1/\lambda_0$, recalling that v_x is bounded, we deduce that $\mathbb{P}_x(X_t = y) \approx u_y$ for all $x, y \in E$.

Proof. The upper bound in (9) follows directly from (8). For the lower bound, we use the upper bound and the fact that $\sum_{x \in E} u_x = \sum_{x \in E} u_x v_x = 1$ to obtain

$$\begin{aligned} u_x v_x &= 1 - \sum_{y \in E, y \neq x} u_y v_y \\ &\geq 1 - e^{\lambda_0 t} \sum_{y \in E, y \neq x} (u_y - C e^{-(\lambda_1 - \lambda_0)t} u_y) \\ &= 1 - e^{\lambda_0 t} + u_x e^{\lambda_0 t} - C e^{-(\lambda_1 - \lambda_0)t} + C u_x e^{-(\lambda_1 - \lambda_0)t}. \end{aligned}$$

The lower bound in (9) follows.

To obtain (10), we combine (6), (8) and (9) in

$$\begin{aligned} &|\mathbb{P}_x(X_t = y) - u_y| \\ &\leq |\mathbb{P}_x(X_t = y \mid t < \tau_\partial) - u_y| \mathbb{P}_x(t < \tau_\partial) + u_y |\mathbb{P}_x(t < \tau_\partial) - v_x e^{-\lambda_0 t}| + u_y |v_x e^{-\lambda_0 t} - 1| \\ &\leq C e^{-(\lambda_1 - \lambda_0)t} (v_x e^{-\lambda_0 t} + C e^{-\lambda_1 t}) + C u_y e^{-\lambda_1 t} + C u_y e^{-\lambda_1 t} + \frac{u_y}{u_x} (1 - e^{-\lambda_0 t} + C e^{-\lambda_1 t}). \end{aligned}$$

We also obtain estimates on the distribution of $(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})$.

Proposition 4.2. *For all $x, y \in E$, all $z \in \partial$ and all $t \geq 0$,*

$$\left| \mathbb{P}_x(X_{\tau_{\partial}^-} = y, X_{\tau_{\partial}} = z) - \frac{u_y q_{y,z}}{\lambda_0} \right| \leq 2C \frac{1 - e^{-\lambda_0 t}}{u_x} + C e^{-\lambda_1 t} (4 + v_x + C e^{-(\lambda_1 - \lambda_0)t}). \quad (11)$$

Again, the inequality (11) gives an accurate estimate on the distribution of $(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})$ in the case where $\lambda_0 \ll \lambda_1$: taking $1/\lambda_1 \ll t \ll 1/\lambda_0$, it follows that $\mathbb{P}_x(X_{\tau_{\partial}^-} = y, X_{\tau_{\partial}} = z) \approx \lambda_0^{-1} u_y q_{y,z}$ for all $x, y \in E$ and $z \in \partial$.

Proof. The proof makes use of the estimates of Proposition 3.1 and the properties of Proposition 3.2: for all bounded measurable function f on $E \times \partial$,

$$\begin{aligned} & \left| \mathbb{E}_x [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}}) \mathbb{1}_{t < \tau_{\partial}}] - v_x e^{-\lambda_0 t} \mathbb{E}_u [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})] \right| \\ & \leq \mathbb{P}_x(t < \tau_{\partial}) \left| \mathbb{E}_x [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}}) \mid t < \tau_{\partial}] - \mathbb{E}_u [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}}) \mid t < \tau_{\partial}] \right| \\ & \quad + \left| \mathbb{E}_u [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})] \right| \left| \mathbb{P}_x(t < \tau_{\partial}) - v_x e^{-\lambda_0 t} \right| \\ & \leq (e^{-\lambda_0 t} v_x + C e^{-\lambda_1 t}) C \|f\|_{\infty} e^{-(\lambda_1 - \lambda_0)t} + C \|f\|_{\infty} e^{-\lambda_1 t}. \end{aligned}$$

This entails

$$\left| \mathbb{E}_x [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}}) \mathbb{1}_{t < \tau_{\partial}}] - v_x e^{-\lambda_0 t} \mathbb{E}_u [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})] \right| \leq C e^{-\lambda_1 t} \|f\|_{\infty} (1 + v_x + C e^{-(\lambda_1 - \lambda_0)t}).$$

Therefore, it follows from the inequality

$$\mathbb{P}_x(\tau_{\partial} \leq t) \leq 1 - v_x e^{-\lambda_0 t} + C e^{-\lambda_1 t}$$

that

$$\begin{aligned} & \left| \mathbb{E}_x [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})] - \mathbb{E}_u [f(X_{\tau_{\partial}^-}, X_{\tau_{\partial}})] \right| \\ & \leq 2 \|f\|_{\infty} (1 - v_x e^{-\lambda_0 t}) + C e^{-\lambda_1 t} \|f\|_{\infty} (2 + v_x + C e^{-(\lambda_1 - \lambda_0)t}). \end{aligned}$$

We conclude from (9) and Proposition 3.2. \square

The previous results are sharp under the condition $\lambda_0 \ll \lambda_1$, which means that absorption takes a long time (the typical absorption time is $1/\lambda_0$) and the process has a tendency to stay away from the absorbing boundaries. This is for example the case when the random walk converges to a deterministic process for which the interior of the domain is stable (see [13, 7, 8]). Our estimates can be applied to the examples of Section 2 by numerically computing the eigenvalues λ_0 and λ_1 .

4.2 The multi-dimensional diffusion model

The previous section gave results on absorbed Markov processes in finite state space, like those of the examples of Section 2. It is also common to model deadlocks replacing discrete random walks with diffusion processes in subsets of \mathbb{R}^d with partly absorbing and reflecting boundaries, like in the examples presented in Section 2. In this case, quasi-stationary may still be defined, although asymptotic properties as those of Propositions 3.1, and hence estimates as in Proposition 4.1, are harder to obtain. A general criterion for such results was recently obtained in [4], which has been applied to various classes of stochastic processes in [3, 4, 5, 6]. A particular case of diffusion in a domain delimited by hyperplanes was studied in [6] using non-linear Lyapunov criteria. However, situations with parts of the boundary being absorbing and other parts reflecting were never studied. We are able to obtain the next result for general colliding stacks models.

Theorem 4.3. *Consider the diffusion process X evolving in $E \cup \partial$ as defined in (3) and (4) with $m_j = 1$ for $j \in \{1, \dots, r\}$ arbitrary nonnegative vectors v_j , $a_i = +\infty$ for $i \in \{1, \dots, d\}$, with the hyperplanes $x_i = 0$, $i \in \{1, \dots, d\}$ as normal reflecting boundaries and the hyperplanes $\langle x, v_j \rangle = 1$ as absorbing. Assume that the infinitesimal generator of X is given for all smooth function f vanishing on the hyperplanes $\langle x, v_j \rangle = 1$ and with zero normal gradient on the hyperplanes $x_i = 0$, by*

$$Lf(x) = \sum_{i,k=1}^d a_{ik}(x) \frac{\partial^2 f}{\partial x_i \partial x_k}(x) + \sum_{i=1}^d b_i(x) \frac{\partial f}{\partial x_i}(x), \quad (12)$$

where the matrix $a = (a_{ij})_{1 \leq i,j \leq d}$ is symmetric and uniformly elliptic and b is uniformly bounded, both Hölder continuous on $E \cup \partial$. Then X admits a unique quasi-stationary distribution α and there exist positive constants $C, \gamma > 0$ such that, for all $t \geq 0$ and all probability measure μ on E ,

$$\|\mathbb{P}_\mu(X_t \in \cdot \mid t < \tau_\partial) - \alpha\|_{TV} \leq Ce^{-\gamma t}. \quad (13)$$

The conclusion of Theorem 4.3 is equivalent to Condition (A) in [4]. This has several implications. For instance, $e^{\lambda_0 t} \mathbb{P}_x(t < \tau_\partial)$ converges when $t \rightarrow +\infty$, uniformly in x , to a positive, bounded eigenfunction η of the infinitesimal generator of $(X_t, t \geq 0)$ for the eigenvalue $-\lambda_0$, characterized by the relation $\mathbb{P}_\alpha(t < \tau_\partial) = e^{-\lambda_0 t}$, $\forall t \geq 0$ [4, Proposition 2.3]. Moreover, it implies a spectral gap property [4, Corollary 2.4], the existence and exponential

ergodicity of the so-called Q -process, defined as the process X conditioned to never hit the absorbing part of the boundary [4, Theorem 3.1] and a conditional ergodic property [5].

Note that the assumption that the constants m_j are all equal to 1 may be relaxed by applying a linear scaling of coordinates, possibly different for each coordinate. Such a scaling has no impact on the required assumptions of the coefficients of the diffusion.

Proof of Theorem 4.3. We start by considering the case where $b_i(x) = 0$ whenever $x_i = 0$. We shall extend, in a second step, our result to the general case.

We extend the definition of σ and b on

$$E' = \{x \in \mathbb{R}^d, \langle |x|, |v_j| \rangle < 1 \forall j \in \{1, \dots, d\}\},$$

where $|x|$ is defined as the vector $(|x_1|, \dots, |x_d|)$, by symmetry over the hyperplanes $x_i = 0$, $i \in \{1, \dots, d\}$: more precisely, for all $x \in E$, we set $a(y) = a(x)$ if $|y| = x$ and $b_i(y) = \text{sign}(y_i)b_i(x)$ for all $i = 1, \dots, d$. Since $b_i(x) = 0$ whenever $x_i = 0$, the extended coefficients are also Hölder continuous on E' . Hence, we can define a diffusion process Y evolving in E' with infinitesimal generator (12), absorbed at the boundary of E' . It follows from standard properties of diffusions with normal reflexion on hyperplanes that the process $(|Y_t^1|, \dots, |Y_t^d|)_{t \geq 0}$ has the same law as the process X defined in the statement of Theorem 4.3. In the proof, we show that (13) holds true for Y , so that it holds true for X .

Since L is assumed to be elliptic, there exist two constants $\bar{\lambda} \geq \underline{\lambda} > 0$ such that

$$\bar{\lambda}\|x\|^2 \geq \langle x, a(x)x \rangle \geq \underline{\lambda}\|x\|^2, \quad \text{for all } x \in E'.$$

In order to prove that Condition (A) of [4] holds true, we use the Lyapunov type criterion proved in [6, Proposition 2.7], with the functions

$$\begin{aligned} \varphi(x) &:= \prod_{\varepsilon \in \{-1, 1\}^d} \prod_{j=1}^r (1 - \langle \varepsilon v_j, x \rangle)^\alpha, \\ V(x) &:= \prod_{\varepsilon \in \{-1, 1\}^d} \prod_{j=1}^r (1 - \langle v_j, x \rangle)^\beta, \end{aligned}$$

where εv_j is intended as a componentwise product $\varepsilon v_j = (\varepsilon_i v_{j,i})_{1 \leq i \leq d}$,

$$\alpha = \frac{1 + \|a\|_\infty \frac{4d^3 2^d r \bar{u}}{u^2}}{\lambda} \quad \text{and} \quad \beta = \frac{\lambda}{2r 2^d \bar{\lambda}} \wedge \frac{1}{2^d r},$$

where $\bar{u} = \sup_{1 \leq j \leq r, 1 \leq i \leq d} v_{j,i}$ and $\underline{u} = \inf_{i,j \text{ s.t. } v_{j,i} > 0} v_{j,i}$. Note that both V and φ are smooth functions on E' .

Assumptions 2 and 4 in [6] are satisfied because L is, in particular, locally elliptic with Hölder coefficients (the arguments are detailed in [6, Section 4]). Assumption 3 in [6] is an easy consequence of the boundedness of E' and of the uniform ellipticity of L . One easily checks by standard arguments (see for instance [3, Section 3]) that, for all $t_0 > 0$, there exist a constant $A > 0$ such that

$$\mathbb{P}_x(t_0 < \tau_\partial) \leq A d(x, \partial E').$$

Since $d(x, \partial E') = \min_{\varepsilon \in \{-1, 1\}^d} (1 - \langle \varepsilon v_j, x \rangle)$, we deduce that

$$\mathbb{P}_x(t_0 < \tau_\partial) \leq A \prod_{\varepsilon \in \{-1, 1\}^d} \prod_{j=1}^r (1 - \langle \varepsilon v_j, x \rangle)^{1/2^d r} \leq A V(x),$$

which is Condition (2.9) of Assumption 1 in [6].

Hence, it remains to prove that there exist a compact set $K \subset E'$ and positive constants $\delta, c, c', c'' > 0$ such that, for all $x \in E'$,

$$L\varphi(x) \geq -c \mathbf{1}_{x \in K} \quad \text{and} \quad LV + c' \frac{V^{1+\delta}}{\varphi^\delta} \leq c'' \varphi. \quad (14)$$

Indeed, by [6, Proposition 2.7], this implies that Assumption 1 in [6] is satisfied, so that (13) holds true (by [6, Corollary 2.8]).

Setting

$$f_i(x) := \sum_{\varepsilon \in \{-1, 1\}^d} \sum_{j=1}^r \frac{\varepsilon_i v_{j,i}}{1 - \langle \varepsilon v_j, x \rangle}, \quad \forall i \in \{1, \dots, d\},$$

we have, for all $x \in E'$,

$$\begin{aligned} L\varphi(x) &= \langle a(x) \nabla, \nabla \rangle \varphi(x) + \langle b(x), \nabla \rangle \varphi(x) \\ &= -\alpha \langle a(x) \nabla, \varphi(x) f(x) \rangle - \alpha \langle b(x), \varphi(x) f(x) \rangle \\ &= \alpha \varphi(x) \langle a(x) (\alpha f(x) - \nabla), f(x) \rangle - \alpha \varphi(x) \langle b(x), f(x) \rangle \\ &\geq \alpha \varphi(x) \left[\alpha \lambda \|f(x)\|^2 - \|a\|_\infty \langle \nabla, f \rangle - \sqrt{d} \|b\|_\infty \|f(x)\| \right]. \end{aligned}$$

Now,

$$\begin{aligned}\langle \nabla, f \rangle(x) &= \sum_{\varepsilon \in \{-1,1\}^d} \sum_{i=1}^d \frac{\|v_j\|^2}{(1 - \langle \varepsilon v_j, x \rangle)^2} \\ &\leq 2^d d r \bar{u} \max_{1 \leq j \leq r} \frac{1}{|1 - \langle v_j, |x| \rangle|^2}.\end{aligned}$$

Let $K_0 = \{x \in E' : \forall j \in \{1, \dots, r\}, \langle v_j, x \rangle \leq 1/2\}$. Our goal is to prove that $\langle \nabla, f \rangle(x) \leq C_0 \|f(x)\|^2$ for all $x \notin K_0$ for an appropriate constant C_0 . We can assume without loss of generality that $x \geq 0$, meaning that all its coordinates are nonnegative. We have for all $1 \leq i \leq d$

$$\begin{aligned}f_i(x) &= \sum_{\varepsilon^{(i)} \in \{-1,1\}^{d-1}} \sum_{j=1}^r v_{ji} \frac{2v_{ji}x_i}{(1 - v_{ji}x_i - \langle \varepsilon^{(i)} u_j^{(i)}, x^{(i)} \rangle)(1 + v_{ji}x_i - \langle \varepsilon^{(i)} u_j^{(i)}, x^{(i)} \rangle)} \\ &\geq \sum_{\varepsilon^{(i)} \in \{-1,1\}^{d-1}} \sum_{j=1}^r \frac{u_{ji}^2 x_i}{1 - v_{ji}x_i - \langle \varepsilon^{(i)} u_j^{(i)}, x^{(i)} \rangle} \\ &\geq \sum_{j=1}^r \frac{u_{ji}^2 x_i}{1 - \langle v_j, x \rangle},\end{aligned}$$

where $x^{(i)}$ denotes the vector of \mathbb{R}^{d-1} obtained from x by suppressing its i -th coordinate. Hence

$$\begin{aligned}\|f(x)\|^2 &\geq \sum_{i=1}^d \left(\sum_{j=1}^r \frac{u_{ji}^2 x_i}{1 - \langle v_j, x \rangle} \right)^2 \\ &\geq \sum_{i=1}^d \sum_{j=1}^r \frac{u_{ji}^4 x_i^2}{|1 - \langle v_j, x \rangle|^2}.\end{aligned}$$

Let j_0 be such that $\langle v_{j_0}, x \rangle$ is maximal. In particular, it is larger than $1/2$, and there exists i_0 such that $v_{j_0 i_0} x_{i_0} \geq 1/2d$. Then,

$$\|f(x)\|^2 \geq \frac{u_{j_0 i_0}^4 x_{i_0}^2}{|1 - \langle v_{j_0}, x \rangle|^2} \geq \frac{\underline{u}^2}{4d^2} \max_{1 \leq j \leq r} \frac{1}{|1 - \langle v_j, |x| \rangle|^2} \geq \frac{\underline{u}^2}{4d^3 2^d r \bar{u}} \langle \nabla, f \rangle(x).$$

So by definition of α

$$L\varphi(x) \geq \alpha\varphi(x)(\|f(x)\|^2 - \sqrt{d}\|b\|_\infty\|f(x)\|).$$

But $\|f\|$ converges to $+\infty$ when $x \rightarrow \partial E'$, so that the left hand side of (14) holds true.

Let us now prove that second part also holds true. We have

$$\begin{aligned} LV(x) &= \beta V(x) \langle a(x)(\beta f(x) - \nabla), f(x) \rangle - \beta V(x) \langle b(x), f(x) \rangle \\ &\leq \beta V(x) [\beta \bar{\lambda} \|f(x)\|^2 - \langle a(x)\nabla, f(x) \rangle + \langle b(x), f(x) \rangle]. \end{aligned}$$

Since

$$\begin{aligned} -\langle a(x)\nabla, f(x) \rangle &= - \sum_{\varepsilon \in \{-1,1\}^d} \sum_{j=1}^r \sum_{i=1}^d \sum_{k=1}^d a_{ik} \frac{\varepsilon_i v_{ji} \varepsilon_k v_{jk}}{(1 - \langle \varepsilon v_j, x \rangle)^2} \\ &\leq -\underline{\lambda} \sum_{\varepsilon \in \{-1,1\}^d} \sum_{j=1}^r \frac{\|v_j\|^2}{(1 - \langle \varepsilon v_j, x \rangle)^2} \\ &= -\underline{\lambda} \langle \nabla, f \rangle(x), \end{aligned}$$

and, by Cauchy-Schwarz inequality,

$$\|f(x)\|^2 \leq r 2^d \langle \nabla, f \rangle(x), \quad (15)$$

we deduce from the definition of β that

$$\begin{aligned} LV(x) &\leq \beta V(x) \left[-\frac{\underline{\lambda}}{2} \langle \nabla, f \rangle(x) + \langle b(x), f(x) \rangle \right] \\ &\leq -\beta V(x) \sum_{\varepsilon \in \{-1,1\}^d} \sum_{j=1}^r \left(\frac{\underline{\lambda} \|v_j\|^2 / 2}{(1 - \langle \varepsilon v_j, x \rangle)^2} - \sum_{i=1}^d \frac{\varepsilon_i v_{ji} b_i(x)}{1 - \langle \varepsilon v_j, x \rangle} \right) \\ &\leq -\frac{\beta \underline{\lambda}}{4} V(x) \sum_{\varepsilon \in \{-1,1\}^d} \sum_{j=1}^r \frac{\|v_j\|^2}{(1 - \langle \varepsilon v_j, x \rangle)^2} + B \\ &\leq -\frac{\beta \underline{\lambda}}{4r 2^d} V(x) \|f(x)\|^2 + B \end{aligned}$$

for some positive constant $B > 0$. Moreover, choosing $\delta = \frac{\beta}{\alpha - \beta}$, one has $\frac{V(x)^{1+\delta}}{\varphi(x)^\delta} = 1$ and hence

$$LV(x) + \frac{V(x)^{1+\delta}}{\varphi(x)^\delta} \leq -\frac{\beta \underline{\lambda}}{4r 2^d} V(x) \|f(x)\|^2 + B + 1,$$

where $-\frac{\beta\lambda}{4r^{2d}}V(x)\|f(x)\|^2 + B + 1$ is non-positive in a vicinity of the boundary $\partial E'$ since $V(x) \geq d(x, \partial E')$ as was proved above. Since φ is uniformly bounded from below by a positive constant on any compact subset of E' and is positive on E' , we deduce that $LV(x) + \frac{V(x)^{1+\varepsilon}}{\varphi(x)^\varepsilon}$ is smaller than $c''\varphi(x)$ for some constant $c'' > 0$. As a consequence, the right hand side of (14) holds true, which concludes the proof when $b_i(x) = 0$ whenever $x_i = 0$.

It only remains to extend the last result to cases where $b_i(x)$ does not vanish when $x_i = 0$. Since by symmetry the functions V and φ satisfy Neumann's boundary condition on the reflecting boundary of E , they both belong to the domain of the generator of the process X , hence we can reproduce the computations above which are actually valid for any bounded measurable b and conclude using the same criterion. This concludes the proof of Theorem 4.3. \square

We can deduce from (13) similar estimates as in Proposition 4.1, with the additional difficulty that we cannot obtain pointwise estimates on η and α , but only estimates on their mean values on small balls. Again, since η is bounded, these estimates are good when $\lambda_0 \ll \lambda_1$.

Proposition 4.4. *For all $t \geq 0$ and $x \in E$,*

$$\eta(x) \leq e^{\lambda_0 t} + C e^{-(\lambda_1 - \lambda_0)t} \quad (16)$$

and, for all measurable $A \subset E$ such that $\alpha(A) > 0$,

$$\frac{1}{\alpha(A)} \int_A \eta(x) d\alpha(x) \geq e^{\lambda_0 t} + C e^{-\gamma t} - \frac{e^{\lambda_0 t} - 1 + C e^{-\gamma t}}{\alpha(A)}. \quad (17)$$

In addition, for all $B \subset E$ measurable such that $\alpha(B) > 0$, we define the probability measure $\alpha|_B$ as $\frac{\mathbb{1}_{x \in B} d\alpha(x)}{\alpha(B)}$. Then, for all such B , all $t \geq 0$ and all $A \subset E$ measurable, we have

$$\begin{aligned} \left| \mathbb{P}_{\alpha|_B}(X_t \in A) - \alpha(A) \right| &\leq \frac{\alpha(A)}{\alpha(B)} (1 - e^{-\lambda_0 t} + C e^{-(\lambda_0 + \gamma)t}) \\ &\quad + C e^{-(\lambda_0 + \gamma)t} (\|\eta\|_\infty + 2\alpha(A) + C e^{-\gamma t}). \end{aligned} \quad (18)$$

Proof. As explained above, the property (13) implies (see [4, Proposition 2.3]) that there exist constants C' and $\gamma' > 0$ such that

$$|e^{\lambda_0 t} \mathbb{P}_x(t < \tau_\partial) - \eta(x)| \leq C' e^{-\gamma' t}, \quad \forall x \in E, \forall t \geq 0. \quad (19)$$

In addition, $\int_E \eta(x) \alpha(dx) = 1$. Without loss of generality, we shall assume that $C' = C$ and $\gamma' = \gamma$ in the inequalities (13) and (19).

We first deduce from (19) the inequality (16). Combining this with $\alpha(E) = \int_E \eta d\alpha = 1$, we deduce that

$$\begin{aligned} \int_A \eta(x) d\alpha(x) &= 1 - \int_{E \setminus A} \eta(x) d\alpha(x) \\ &\geq 1 - (e^{\lambda_0 t} + C e^{-\gamma t}) \alpha(E \setminus A) \\ &\geq 1 - e^{\lambda_0 t} - C e^{-\gamma t} + \alpha(A) (e^{\lambda_0 t} + C e^{-\gamma t}). \end{aligned}$$

This entails (17).

To obtain (18), we combine (13), (19) and (17) in

$$\begin{aligned} \left| \mathbb{P}_{\alpha|_B}(X_t \in A) - \alpha(A) \right| &\leq \left| \mathbb{P}_{\alpha|_B}(X_t \in A \mid t < \tau_\partial) - \alpha(A) \right| \mathbb{P}_{\alpha|_B}(t < \tau_\partial) \\ &\quad + \alpha(A) \left| \mathbb{P}_{\alpha|_B}(t < \tau_\partial) - \frac{1}{\alpha(B)} \int_B \eta(x) d\alpha(x) e^{-\lambda_0 t} \right| \\ &\quad + \alpha(A) \left| \frac{1}{\alpha(B)} \int_B \eta(x) d\alpha(x) e^{-\lambda_0 t} - 1 \right| \\ &\leq C e^{-\gamma t} \left(\frac{1}{\alpha(B)} \int_B \eta(x) d\alpha(x) e^{-\lambda_0 t} + C e^{-(\lambda_0 + \gamma)t} \right) \\ &\quad + 2C \alpha(A) e^{-(\lambda_0 + \gamma)t} + \frac{\alpha(A)}{\alpha(B)} (1 - e^{-\lambda_0 t} + C e^{-(\lambda_0 + \gamma)t}). \quad \square \end{aligned}$$

5 Simulations

In this section, we focus on the examples presented in Section 2.

As illustration of the different models studied above, we present simulation results of the density of QSD as well as the value of λ_0 for a diffusion model of two colliding stacks and a discrete state-space model of a 2-dimensional banker algorithm. The numerical method applies to general models. We restrict to the case of two stacks and two consumers for pedagogical reasons.

We use a specific class of particle systems with singular interaction, called Fleming-Viot particle systems, which arises in the study of distributions of absorbed Markov processes conditioned to non-absorption [1, 12]. In these

systems, particles move independently following the Markovian dynamics of the underlying process (here, colliding stacks or banker algorithms), until one gets absorbed, in which case it is immediately sent to the position of another particle, chosen uniformly at random. This method overcomes the problematic and necessary deterioration of classical Monte-Carlo techniques in the setting of absorbed processes, by maintaining a constant sample size of significant particles.

It is known in general that this method allows to approximate conditional distributions of the underlying Markov process in the limit of infinitely many particles [24, 22]. In our case, since we want to approximate the quasi-stationary distribution, we need to simulate the Fleming-Viot process for sufficiently long time. In practice, we compute the ergodic mean of the simulated system and stop the simulation when its variation goes below a threshold.

In this setting, the eigenvalue λ_0 is obtained as the average rate of absorption of particles. This approximation relies on the unbiased estimator introduced in the proof of Theorem 2.1 of [23].


5.1 Simulations for two colliding stacks

We simulate a diffusive model of two colliding stacks, given by the solution $X_t = (X_t^{(1)}, X_t^{(2)})$ of the stochastic differential equation


$$\begin{aligned} dX_t^{(1)} &= dB_t^{(1)} + \sigma dB_t^{(2)} \\ dX_t^{(2)} &= \sigma dB_t^{(1)} + dB_t^{(2)}, \end{aligned}$$

with a parameter $\sigma \geq 0$ and two independent Brownian motions $B^{(1)}$ and $B^{(2)}$. The process is assumed to be killed when $X_t^{(1)} + X_t^{(2)} = 1$ and reflected when $X_t^{(1)} = 0$ or $X_t^{(2)} = 0$. The case $\sigma = 0$ corresponds to two independent Brownian motions for each coordinate. The parameter σ governs the correlation between the two coordinates.


The numerical results are presented in Fig. 1, where the density of the QSD is plotted for different values of σ . We observe that small values of σ have little influence on the density of the QSD (Fig. 1(a) and (b)). Larger values of σ have a tendency to concentrate the density close to the line $x = y$. This is due to the stronger correlation between the two coordinates (with the limiting case $\sigma = 1$ where $X^{(1)} = X^{(2)}$). In all simulations, the density of the QSD vanishes at the absorbing boundaries of the domain, as expected, and




(a) Simulation with $\sigma = 0$. We obtain $\lambda_0 \approx 4.65$.



(b) Simulation with $\sigma = 0.5$. We obtain $\lambda_0 \approx 5.69$.



(c) Simulation with $\sigma = 0.7$. We obtain $\lambda_0 \approx 6.73$.



(d) Simulation with $\sigma = 0.9$. We obtain $\lambda_0 \approx 8.06$.

Figure 1: Density of the QSD for the SDE $dX_t = SdB_t$ in \mathbb{R}^2 , with $S = (\sigma, 1; 1, \sigma)$, killed when $X_t^{(1)} + X_t^{(2)} = 1$ and reflected when $X_t^{(1)} = 0$ or $X_t^{(2)} = 0$.

decreases with respect to $X^{(1)}$ and $X^{(2)}$. Hence the larger density is obtained at the point $(0, 0)$, farthest from the absorbing boundary.

The computed eigenvalues λ_0 increase with σ . This means that the deadlock of the system is faster for larger values of σ . Although the densities of the QSD for $\sigma = 0$ and $\sigma = 0.5$ are quite similar, the values of λ_0 are significantly different.

5.2 Simulations for the banker algorithm in the case of two consumers


We consider here a Markov chain on \mathbb{N}^2 , which is a random walk $X_n = (X_n^{(1)}, X_n^{(2)})$ in discrete time, with the following probabilities of transitions. The process jumps from (i, j) to

$$\begin{aligned} & (i, j) \text{ with probability } \frac{1}{3} \\ & (i + 1, j + 1) \text{ with probability } \frac{1}{6} \\ & (i + 1, j - 1) \text{ with probability } \frac{1}{6} \\ & (i - 1, j + 1) \text{ with probability } \frac{1}{6} \\ & (i - 1, j - 1) \text{ with probability } \frac{1}{6}, \end{aligned}$$


with additional killing of the process when it reaches $X_n^{(1)} + X_n^{(2)} = 100$ and with various schemes of reflection. In all cases, the first (resp. second) coordinate of the process is reflected when $X_n^{(1)} = 0$ (resp. $X_n^{(2)} = 0$). In addition to this, reflection occurs when $X_n^{(1)}$ or $X_n^{(2)}$ reach positive thresholds m_1 and m_2 which vary in the following simulations. Note that if $m_1 \geq 100$ and $m_2 \geq 100$, we are back to a model of two colliding stacks. We added a possibility for the process to remain at the same position to avoid periodicity problems.

The numerical results are presented in Fig. 2, where the QSD is plotted for different values of m_1 and m_2 . This corresponds to various shapes of the domain. We observe again that QSD vanishes at the absorbing boundary and increases with the distance to this boundary. In the three simulations, we observe positive values of the QSD at reflection boundaries.

The computed eigenvalues λ_0 depend both on the shape of the domain of the process (the larger the domain is, the smaller λ_0 should be) and the size of the absorbing part of the boundary (the larger it is, the larger λ_0 should be). Hence, the dependence of λ_0 with respect to m_1 and m_2 is non-trivial. For example, when comparing Fig. 2 (a) and (c), we see that the domain is larger but also the absorbing boundary. In this case, this produces a larger value of λ_0 , hence a higher speed of deadlock, although the amount of resources is larger.



(a) Reflection when $X_1 \in \{0, 80\}$ or $X_2 \in \{0, 80\}$. We obtain $\lambda_0 \approx 2.88 \times 10^{-4}$.
(b) Reflection when $X_1 \in \{0, 70\}$ or $X_2 \in \{0, 90\}$. We obtain $\lambda_0 \approx 2.67 \times 10^{-4}$.



(c) Reflection when $X_1 \in \{0, 80\}$ or $X_2 \in \{0, 100\}$. We obtain $\lambda_0 \approx 2.99 \times 10^{-4}$.

Figure 2: Density of the QSD for a symmetric random walk on \mathbb{N}^2 , killed when $X_n^{(1)} + X_n^{(2)} = 100$ with different schemes of reflection.

6 Concluding remarks and further research aspects

So far as we know, this paper is the first attempt to use the theory of quasi-invariant distributions (QSDs) for analysing distributed algorithms which involve large numbers of processors and resources (i.e. types of memories). Our results are both non-asymptotic (Section 4.1) and asymptotic (Section 4.2)

in time, and non-asymptotic in the number of resources and processors: they are true for any finite values of the involved parameters and reflect the real behavior of these algorithms.

Analysing other types of dynamic algorithms with similar QSD-tools will be the object of further research.

References

- [1] K. Burdzy, R. Holyst and . March. A FlemingViot Particle Representation of the Dirichlet Laplacian. *Communications in Mathematical Physics* **214(3)**, 679-703, 2000.
- [2] P. Cattiaux and S. Méléard. Competitive or weak cooperative stochastic Lotka-Volterra systems conditioned on non-extinction. *J. Math. Biol.* **60(6)**, 797–829, 2010.
- [3] N. Champagnat, A.K. Coulibaly-Pasquier and D. Villemonais. Exponential convergence to quasi-stationary distribution for multi-dimensional diffusion processes. *Arxiv e-prints*, 2017.
- [4] N. Champagnat and D. Villemonais. Exponential convergence to quasi-stationary distribution and Q-process. *Probab. Theory Related Fields*, **164(1)**, 243–283, 2016.
- [5] N. Champagnat and D. Villemonais. Uniform convergence to the Q-process. *Arxiv e-prints*, 2017.
- [6] N. Champagnat and D. Villemonais. Lyapunov criteria for uniform convergence of conditional distributions of absorbed Markov processes. *Arxiv e-prints*, 2017.
- [7] Chazottes, J.-R., Collet, P. and Méléard, S. Sharp asymptotics for the quasi-stationary distribution of birth-and-death processes. *Probab. Theory Related Fields*, **164(1–2)**, 285332, 2016).
- [8] Chazottes, J.-R., Collet, P. and Méléard, S. On time scales and quasi-stationary distributions for multitype birth-and-death processes. Eprint arXiv:1702.05369, 2017.

- [9] F. Comets, F. Delarue and R. Schott. Distributed algorithms in an ergodic markovian environment. *Random Structures and Algorithms*, **30(1–2)**, 131–167, 2007.
- [10] F. Comets, F. Delarue and R. Schott. Large deviations analysis for distributed algorithms in an ergodic markovian environment. *Applied Mathematics and Optimization*, **3–3**, 341–396, 2009.
- [11] J.N. Darroch and E. Seneta. On quasi-stationary distributions in absorbing continuous-time finite Markov chains. *J. Appl. Probability*, **4**, 192–196, 1967.
- [12] Del Moral, Pierre. *Mean field simulation for Monte Carlo integration*. CRC press, 2013.
- [13] M. Faure and S.J. Schreiber. Quasi-stationary distributions for randomly perturbed dynamical systems. *Ann. Appl. Probab.*, **24(2)**, 553–598, 2014.
- [14] P. Flajolet. The evolution of two stacks in bounded space and random walks in a triangle. *Proceedings of MFCS’86*, Lecture Notes in Computer Science 233, 325-340, 1986.
- [15] N. Guillotin-Plantard and R. Schott. Distributed algorithms with dynamic random transitions. *Random Structures and Algorithms*, **21(3–4)**, 371–396, 2002.
- [16] D.E. Knuth. *The art of computer programming*, vol. 1, Addison Wesley, 1973.
- [17] G. Louchard, R. Schott, M. Tolley and P. Zimmermann. Random walks, heat equation and distributed algorithms. *Journal of Computational and Applied Mathematics*. **53**, 243–274, 1994.
- [18] G. Louchard and R. Schott. Probabilistic analysis of some distributed algorithms. *Random Structures and Algorithms*, **2**, 151–186, 1991.
- [19] R.S. Maier. Colliding stacks: a large deviations analysis. *Random Structures and Algorithms*, **2**, 379-420, 1991b.

- [20] R.S. Maier and R. Schott. Exhaustion of shared memory: Stochastic results. *Proceedings of WADS'93*, LNCS 709, 494-505, Springer Verlag, 1993.
- [21] S. Méléard and D. Villemonais. Quasi-stationary distributions and population processes. *Probab. Surv.*, **9**, 340–410, 2012.
- [22] W. Oçafrain and D. Villemonais. Convergence of a non-failable mean-field particle system. *Stochastic Analysis and Applications*, 35(4), 587-603, 2017.
- [23] D. Villemonais. General approximation method for the distribution of Markov processes conditioned not to be killed. *ESAIM: Probability and Statistics*, **18**, 441-467, 2014.
- [24] D. Villemonais. *Distributions quasi-stationnaires et mthodes particulaires pour l'approximation de processus conditionns*. Doctoral dissertation, Ecole Polytechnique X, 2011.
- [25] A.C. Yao. An analysis of a memory allocation scheme for implementing stacks. *SIAM J. Comput.*, **2**, 398-403, 1981.