



An approximation scheme for quasi-stationary distributions of killed diffusions

Andi Q. Wang^{a,*}, Gareth O. Roberts^b, David Steinsaltz^a

^a Department of Statistics, University of Oxford, 24–29 St Giles', Oxford, OX1 3LB, UK

^b Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK

Received 21 August 2018; received in revised form 1 July 2019; accepted 18 September 2019

Available online 23 September 2019

Abstract

In this paper we study the asymptotic behavior of the normalized weighted empirical occupation measures of a diffusion process on a compact manifold which is killed at a smooth rate and then regenerated at a random location, distributed according to the weighted empirical occupation measure. We show that the weighted occupation measures almost surely comprise an asymptotic pseudo-trajectory for a certain deterministic measure-valued semiflow, after suitably rescaling the time, and that with probability one they converge to the quasi-stationary distribution of the killed diffusion. These results provide theoretical justification for a scalable quasi-stationary Monte Carlo method for sampling from Bayesian posterior distributions.

© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

MSC: primary 60B12; 60J60; 37C50; secondary 65C05

Keywords: Asymptotic pseudo-trajectory; Killed diffusion; Quasi-stationary distribution; Quasi-stationary Monte Carlo method; Stochastic approximation

1. Introduction

In this paper we consider a stochastic approximation algorithm to approximate the quasi-stationary distribution of a killed diffusion on a compact manifold. This is particularly motivated by recent developments in theoretical and computational statistics: a promising new paradigm has emerged for the performance of exact Bayesian inference on large datasets, known as *quasi-stationary Monte Carlo* (QSMC) methods, [23,28]. Recall that for a Markov process $Y = (Y_t)_{t \geq 0}$ killed at random time τ_∂ , a distribution π is called *quasi-stationary* if for

* Corresponding author.

E-mail addresses: a.wang@stats.ox.ac.uk (A.Q. Wang), gareth.o.roberts@warwick.ac.uk (G.O. Roberts), steinsal@stats.ox.ac.uk (D. Steinsaltz).

<https://doi.org/10.1016/j.spa.2019.09.010>

0304-4149/© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

each $t \geq 0$,

$$\mathbb{P}_\pi(Y_t \in \cdot | \tau_\partial > t) = \pi(\cdot) \quad (1)$$

where $\mathbb{P}_\pi = \int \pi(dy) \mathbb{P}_y$, \mathbb{P}_y denoting the law of Y with initial position $Y_0 = y$. The idea behind quasi-stationary Monte Carlo is to draw samples from such a quasi-stationary distribution.

There are (at least) two distinct approaches to realizing this method in practice. The first is to use particle approximation methods. This was the approach of the original QSMC method, the Scalable Langevin Exact (ScaLE) algorithm, which was introduced in [23]. The quasi-stationary framework enables the principled use of *subsampling* techniques, and the resulting ScaLE algorithm is provably efficient for performing exact Bayesian inference when the underlying dataset is large. This approach is supported by the full weight of the sequential Monte Carlo literature, and thus has theoretical assurances of convergence.

While the ScaLE algorithm possesses some desirable theoretical and computational properties, it also suffers several drawbacks. One particularly prominent drawback is the high complexity of implementing the algorithm. A cursory glance through [23] can convince the reader that the algorithm is (necessarily) involved. This motivates the search for alternative QSMC methods which inherit the desirable properties of exactness and scalability, while avoiding the algorithmically complex particle approximation framework.

One alternative approach to QSMC takes advantage of the fact that quasi-stationary distributions can be written as solutions to fixed point equations in measure spaces (represented here by Eq. (18) in Proposition 2.12). This yields the QSMC method dubbed Regenerating ScaLE (ReScaLE), which seeks to approximate the quasi-stationary distribution π by means of stochastic approximation. The idea behind ReScaLE is that we simulate a *single* killed diffusion path, whose quasi-stationary distribution coincides with the Bayesian posterior of interest. When this single particle is killed, the particle is ‘regenerated’ or ‘reborn’ in a new random location drawn independently from its normalized weighted empirical measure. It then continues to evolve from this new starting position until it is killed again, at which point it is again regenerated, and so on. This algorithm is inspired by similar algorithms for the approximation of quasi-stationary distributions in discrete-time settings [1,3,4,8]. Indeed, even outside of QSMC such algorithms have been applied to the simulation of quasi-stationary distributions in ecology [26].

The resulting ReScaLE algorithm inherits the key properties that motivate the ScaLE algorithm, and algorithmically it is significantly more transparent and straightforward, but this is worthwhile only if the method provably leads to correct results. Theoretical analysis of stochastic approximation approaches to numerically solving fixed point equations are well-studied in finite-dimensional contexts, *cf.* [21], but there is currently no theory appropriate for the measure-valued, continuous-time context of ReScaLE. While the implementation of the original ScaLE algorithm is underpinned by the vast body of work on sequential Monte Carlo, there has been no theoretical assurance that the ReScaLE algorithm even converges. The purpose of the present work is to demonstrate this fundamental convergence property, leaving practical and computational properties of the algorithm to be considered in detail in forthcoming work.

In this work we assume that we are working on a compact, boundaryless, connected, d -dimensional smooth Riemannian manifold M , as in [5]. Relaxing the assumption of compactness is discussed briefly in Section 5. We have a particle $X = (X_t)_{t \geq 0}$ evolving on M in continuous time, according to the solution to the stochastic differential equation (SDE)

$$dY_t = \nabla A(Y_t)dt + dW_t \quad (2)$$

between regeneration events, where $A : M \rightarrow \mathbb{R}$ is a smooth function and W is a standard Brownian motion on M . Regeneration events occur at a state-dependent rate $\kappa(X_{t-})$, where $\kappa : M \rightarrow [0, \infty)$ is a given smooth, positive function, which we will refer to as the *killing rate*. At a regeneration event, say at time T , the particle is instantaneously ‘killed’ and ‘reborn’: its location is drawn (independently) from its normalized weighted empirical occupation measure μ_T , where for all $t \geq 0$, μ_t is given by

$$\mu_t = \frac{r\mu_0}{r + \int_0^t \eta_s ds} + \frac{\int_0^t \eta_s \delta_{X_s} ds}{r + \int_0^t \eta_s ds}$$

where $r > 0$, $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and a probability measure μ_0 on M are fixed. The resulting process X is clearly non-Markovian. The addition of the μ_0 term has the benefit of regularizing the μ_t around $t = 0$, as well as providing practical flexibility for the resulting Monte Carlo algorithm. The weight function η similarly provides additional practical flexibility; a straightforward choice would be constant $\eta_s \equiv 1$, but see Remark 2.2 for a nonconstant alternative.

1.1. Main results

The goal of this paper is to characterize the asymptotic behavior of the measure-valued process $(\mu_t)_{t \geq 0}$, and to show that it converges to the quasi-stationary distribution for the original killed process. We proceed in two steps, following the “ODE method” (cf. [2,21]), which has been used to prove convergence of similar reinforced processes, for instance in [3–5,20]. The ODE method proceeds by two key steps. First, showing that a certain deterministic semiflow Φ converges to the appropriate limit, in our case the quasi-stationary distribution π . Second, showing that, following a suitable deterministic time change $\zeta_t := \mu_{h(t)}$, the stochastic evolution of the measures $(\zeta_t)_{t \geq 0}$ shadows Φ in an appropriate sense, to be defined below. From these two properties almost-sure convergence of μ_t to the quasi-stationary distribution can be deduced.

The present paper extends previous related work in considering a continuous-time diffusive process on a compact manifold (rather than, say, a discrete-time Markov chain), which experiences ‘soft killing’ according to a smooth state-dependent killing rate (rather than instantaneous ‘hard killing’ at a boundary). In this setting of soft killing, a generic killing time τ_∂ has the form

$$\tau_\partial = \inf \left\{ t \geq 0 : \int_0^t \kappa(X_s) ds \geq \xi \right\}, \quad (3)$$

where $\xi \sim \text{Exp}(1)$ is independent of the process X , rather than the first hitting time of some forbidden set of states as in the hard killing case.

Theorem 1.1. Under Assumptions 1–4 detailed in Section 2.1, we have with probability 1, that for each $T > 0$,

$$\lim_{t \rightarrow \infty} \sup_{s \in [0, T]} d_w(\zeta_{t+s}, \Phi_s(\zeta_t)) = 0 \quad (4)$$

where d_w is a metric that metrizes weak-* convergence of probability measures given in (5) and Φ is the semiflow defined in Section 3.

Remark 1.2. In [2], a map $t \mapsto \zeta_t$ that satisfies (4) for each $T > 0$ is called an *asymptotic pseudo-trajectory* for the measure-valued semiflow Φ .

In Section 3 we define the semiflow Φ , and proceed to show that it has a global attractor π , which is the unique quasi-stationary distribution of the diffusion (2) killed at rate κ .

In particular, Theorem 1.1 leads to the following corollary:

Corollary 1.3. *Under the conditions of Theorem 1.1, we have that almost surely, $\lim_{t \rightarrow \infty} \mu_t = \pi$ in the sense of weak-* convergence.*

That is, we have $\lim_{t \rightarrow \infty} \mu_t(f) = \pi(f)$ for any continuous $f : M \rightarrow \mathbb{R}$.

While our results above hold for any appropriate given killing rate which satisfies our assumptions, in the ReScaLE algorithm the killing rate κ is chosen so that the quasi-stationary distribution π equals the Bayesian posterior distribution of interest; see the expression (7), cf. [28]. Corollary 1.3 tells us that in this setting we can draw approximate samples from π by running the regenerating process X and outputting μ_t for a large t as a proxy for π .

Fig. 1 shows the output of two independent simulations on the unit circle $M = \mathbb{R}/(2\pi\mathbb{Z})$ parameterized by $\theta \in [0, 2\pi)$, which is amenable to straightforward simulation and visualization. The underlying diffusion is a Brownian motion ($A \equiv 0$ in (2)), and the quasi-stationary distribution is trimodal with density function $\pi(\theta) = (0.3 + \sin^2(1.5\theta))/(1.6\pi)$ with respect to the Riemannian measure. The corresponding killing rate, calculated using (7) is $\kappa(\theta) = 2.25(2\cos^2(1.5\theta) - 1)/(0.3 + \sin^2(1.5\theta)) + K$, with $K = 1.75$.

The simulations were done using a simple Euler discretization scheme where time was discretized into intervals of length 0.05. The plots are oriented so that $\theta = 0$ is due east, with $\theta = \pi/2$ being due north, and so on. We have chosen $r = 1000$, μ_0 to be the uniform distribution over M and $\eta_s \equiv 1$. The plots show $\int_0^t \delta_{X_s} - ds/t$ for $t = 25, 100, 1000, 10^6$, split into 50 evenly-spaced bins. The quasi-stationary distribution π is the dashed line.

We see that there is a significant amount of variability for the small time values, as the initial rebirths are drawn mostly from the uniform distribution μ_0 . These discrepancies are largely smoothed out by $t = 1000$ and certainly by $t = 10^6$, $\int_0^t \delta_{X_s} - ds/t$ closely resembles the quasi-stationary distribution π .

This present work closely follows the approaches of [5] and the recent [4] in order to prove our limiting result Theorem 1.1. In [5] the authors show an analogous characterization for the normalized occupation measures of a self-interacting diffusion on a compact space, where the empirical occupation measure influences the present behavior of the diffusion through its drift term. In our work, the influence of the occupation measures is felt through jumps, which occur at a state-dependent rate. In [4], the authors prove a discrete-time analogue of our above result; their underlying Markov process is a discrete-time Markov chain evolving on a compact space, rather than a diffusion. In [4, Section 8.3], the authors do suggest a continuous-time extension of their work: a diffusive process that is killed instantaneously when hitting the boundary of an open set. This is not the same as the present work; we are assuming a boundaryless manifold, and instead of hard killing at a boundary, killing occurs at a smooth state-dependent rate κ as in (3). The key difference in the proposed setting of [4, Section 8.3] is that in their case, the state space is no longer compact, and hence additional arguments ensuring almost sure tightness of the empirical occupation measures are needed.

We follow generally the path mapped out by [4,5,20], often referred to as the “ODE method”, cf. [2,21]. To understand the particular continuous-time dynamics we employ techniques similar to those of [19,28]. For example, to handle the killing mechanism we make use of the transition subdensity of the killed diffusion and the corresponding resolvent operator; see Lemmas 2.4, 2.5 and 2.8. The continuous-time setting also provides a natural interpretation of the weights η_s in terms of the distribution of the rebirth ‘times’; see Remark 2.2. Another contribution of




Fig. 1. Example on the circle. The circle is parameterized by $\theta \in [0, 2\pi)$, with due east corresponding to $\theta = 0$. The underlying diffusion is a Brownian motion, and we have chosen $r = 1000$, μ_0 to be the uniform distribution on M and $\eta_s \equiv 1$. The quasi-stationary distribution $\pi(\theta) = (0.3 + \sin^2(1.5\theta))/(1.6\pi)$ is the dashed line. We have plotted $\int_0^t \delta_{X_s} - ds/t$ discretized into 50 evenly-spaced bins, for $t = 25, 100, 1000, 10^6$. The two columns are two independent runs.

this work is the analysis of the deterministic semiflow in Section 3, where we transform to an auxiliary Markov process and apply an appropriate drift condition.

The earliest version of a similar analysis that we are aware of is the convergence proof of [1] for the finite-state-space discrete-time setting. [3] and [8] expanded this work to derive rates

of convergence and central limit theorems. These analyses rely upon special techniques, such as those in [21], applicable only to finite-dimensional probability distributions.

The structure of the paper is as follows. In Section 2 we lay out the mathematical setting and describe the assumptions that we are making. We also define the fixed rebirth process and prove some of its key properties which are crucial for later defining the semiflow. In Section 3 we define and analyze the deterministic semiflow Φ and prove that π is a global attractor. In Section 4 we prove that our normalized weighted occupation measures almost surely comprise an asymptotic pseudo-trajectory for Φ , concluding the convergence proof. Finally, in Section 5, we briefly discuss the possibilities for extending these results to noncompact manifolds.

2. Preliminaries

2.1. General background

We first describe our assumptions and some notation. These assumptions are *assumed to hold throughout the paper*.

Assumption 1. M is a d -dimensional boundaryless C^∞ compact connected Riemannian manifold.

We will denote the corresponding Riemannian measure by dx or dy , and the Riemannian inner product of points x and y by $x \cdot y$. Let $C(M)$ denote the Banach space of (bounded) real-valued continuous functions on M equipped with the sup norm, $\|\cdot\|_\infty$. Let $\mathcal{P}(M)$ denote the space of Borel probability measures on M , equipped with the topology of weak-* convergence, namely convergence along all bounded continuous test functions. (This is conventionally called simply weak convergence in probability theory, but we follow the terminology of [4] to avoid confusion when working with the Banach space $C(M)$ and its dual.) Thus $\mathcal{P}(M)$ is a compact metrizable space, by the Prokhorov theorem, cf. [15, Theorem 11.5.4].

As M is compact, $C(M)$ is separable, so we may choose a sequence of smooth functions f_1, f_2, \dots which are dense in $\{f \in C(M) : \|f\|_\infty \leq 1\}$. We define the metric $d_w : \mathcal{P}(M) \times \mathcal{P}(M) \rightarrow \mathbb{R}_+$ on $\mathcal{P}(M)$ by

$$d_w(\nu_1, \nu_2) = \sum_{i=1}^{\infty} \frac{1}{2^i} |\nu_1(f_i) - \nu_2(f_i)|. \quad (5)$$

We also define the total-variation norm for a signed measure μ on M by

$$\|\mu\|_1 = \sup\{|\mu(f)| : f \in C(M), \|f\|_\infty \leq 1\}.$$

Let Ω be the Skorokhod space of càdlàg paths $\omega : \mathbb{R}_+ \rightarrow M$, and let \mathcal{F} be the cylinder σ -algebra. Let $X = (X_t)_{t \geq 0}$ be the coordinate process, $X_t(\omega) = \omega(t)$, and let $(\mathcal{F}_t)_{t \geq 0}$ be the natural filtration of X .

Assumption 2. The function $A : M \rightarrow \mathbb{R}$ is C^∞ .

In particular, this implies that A is bounded (since M is compact), and that the SDE (2) on M has a unique strong solution for any initial position.

Recall that for a Markov process Y killed at a random stopping time τ_∂ , a distribution $\pi \in \mathcal{P}(M)$ is *quasi-stationary* if it satisfies (1) for all $t \geq 0$.

Given a measurable killing rate $\kappa : M \rightarrow [0, \infty)$, the corresponding killing time τ_∂ is defined as

$$\tau_\partial := \inf \left\{ t \geq 0 : \int_0^t \kappa(Y_s) ds \geq \xi \right\},$$

where ξ is an independent exponential random variable with rate 1.

Assumption 3. The killing rate $\kappa : M \rightarrow [0, \infty)$ is C^∞ and is uniformly bounded away from zero: there exists some constant $\underline{\kappa} > 0$ such that

$$0 < \underline{\kappa} \leq \kappa(x) \quad \forall x \in M. \quad (6)$$

As a continuous function on a compact space κ is necessarily bounded above, say

$$\kappa(x) \leq \bar{\kappa} < \infty \quad \forall x \in M.$$

Given a killing rate κ which is not strictly bounded away from 0, we can always add a positive constant everywhere; this will not affect the quasi-stationary behavior of the process and will ensure that (6) holds. The upper bound on the killing rate will certainly guarantee that the resulting process will be almost surely nonexplosive.

We will later see that given κ , the diffusion (2) killed at rate κ has a unique quasi-stationary distribution (Proposition 2.12).

The question of existence and uniqueness of quasi-stationary distributions of killed diffusions has been studied in depth, for instance in [19], [11, Chapter 6]. This context of quasi-stationary Monte Carlo methods — where we are starting from a density π and wish to construct a killed diffusion whose quasi-stationary distribution coincides with π — is the topic of [28], and the interested reader is encouraged to look therein for more details.

From [28], if we were to start with a smooth positive density π , the appropriate choice of κ is given for each $y \in M$ by

$$\kappa(y) := \frac{1}{2} \left(\frac{\Delta \pi}{\pi} - \frac{2 \nabla A \cdot \nabla \pi}{\pi} - 2 \Delta A \right)(y) + K. \quad (7)$$

Here ∇ and Δ are the gradient and Laplacian operators on M , and K is a constant chosen so that κ satisfies (6). This choice of κ is discussed at length in [28]. In this work we will not necessarily assume we are starting from π , and take κ to be a general killing rate satisfying Assumption 3.

Fix $\mu_0 \in \mathcal{P}(M)$ and $r > 0$. We fix a weight function $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying the following assumptions: Define functions $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and $\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by

$$g(t) := \int_0^t \eta_s ds, \quad \alpha_t := \frac{\eta_t}{r + g(t)}, \quad t \geq 0.$$

Assumption 4. η is continuously differentiable with $\eta_t > 0$ for all $t > 0$, and $g(t) \rightarrow \infty$ as $t \rightarrow \infty$. α is differentiable and satisfies $\alpha_t \rightarrow 0$ as $t \rightarrow \infty$, $\int_0^\infty \alpha_s^2 ds < \infty$ and

$$\sum_{n=1}^\infty \int_{h(n)}^\infty \alpha_s^2 ds < \infty, \quad (8)$$

where h is defined below.

Since g is strictly increasing (as $\eta_t > 0$), continuously differentiable and increases to ∞ , it is a diffeomorphism of $\mathbb{R}_+ \rightarrow \mathbb{R}_+$. Thus it has a well-defined continuously differentiable inverse g^{-1} . The function $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is then given by

$$h(t) := g^{-1}(re^t - r), \quad t \geq 0.$$

This function h will be the time change that we shall employ in Section 4.

Remark 2.1. It follows from [Assumption 4](#) that $\int_0^t \alpha_s ds = \log(1 + g(t)/r) \rightarrow \infty$ as $t \rightarrow \infty$. Thus these conditions on α are analogous to the typical discrete-time assumptions on the step sizes in traditional stochastic approximation; cf. [21, Chapter 5]. Since $t \mapsto \int_{h(t)}^\infty \alpha_s^2 ds$ is monotone decreasing, a sufficient condition for (8) to hold is that

$$\int_0^\infty dt \int_{h(t)}^\infty ds \alpha_s^2 = \int_0^\infty \alpha_s^2 \log(1 + g(s)/r) ds < \infty.$$

Define the normalized empirical occupation measures $(\mu_t(\omega))_{t \geq 0}$ by

$$\mu_t(\omega) = \frac{r\mu_0 + \int_0^t \eta_s \delta_{\omega(s-)} ds}{r + \int_0^t \eta_s ds}$$

where $\int_0^t \eta_s \delta_{\omega(s-)} ds(A) = \int_0^t \eta_s 1_A(\omega(s-)) ds$ for each measurable $A \subset M$. In general we will omit the dependence on ω .

Remark 2.2. For simplicity one may take $\eta_t \equiv 1$, then $g(t) = t$ and $\alpha_t = 1/(r+t)$ for all $t \geq 0$. Sampling from $Z \sim \int_0^t \eta_s \delta_{X_{s-}} ds / \int_0^t \eta_s ds$ is then equivalent to simulating $V \sim \text{Unif}([0, 1])$ and setting $Z = X_{(Vt)-}$. More generally, for $k \geq 0$ one can take $\eta_t = t^k$ for each $t \geq 0$. It is not difficult to check that for this choice of η [Assumption 4](#) is satisfied, and simulating $Z \sim \int_0^t \eta_s \delta_{X_{s-}} ds / \int_0^t \eta_s ds$ is equivalent to the case of constant η except with $V \sim \text{Beta}(k+1, 1)$. Heuristically, choosing a larger value of k prioritizes the more recent times. The choice of the parameter k to accelerate convergence is itself an interesting question, which will be explored in future work. Preliminary simulations involving Brownian motions and unimodal targets seem to suggest $k = 10$ might be a reasonable choice.

For any $\mu \in \mathcal{P}(M)$, define the operator $-L_\mu$ on smooth C^∞ functions by

$$-L_\mu f(x) = \frac{1}{2} \Delta f(x) + \nabla A \cdot \nabla f(x) + \kappa(x) \int_M (f(y) - f(x)) \mu(dy). \quad (9)$$

Here we choose to use the negative operator in order to comport with the convention adopted in [19,28], where it was chosen to make the corresponding self-adjoint operators positive.

We can define probability measures $(\mathbb{P}_x : x \in M)$ with the following properties:

- $\mathbb{P}_x(X_0 = x) = 1$.
- For all smooth $f \in C(M)$

$$N_t^f := f(X_t) - f(x) - \int_0^t (-L_{\mu_s})f(X_{s-}) ds$$

is a \mathbb{P}_x -martingale with respect to (\mathcal{F}_t) .

We can construct \mathbb{P}_x by explicitly defining the killing events, as follows. Let ξ_1, ξ_2, \dots be an i.i.d. sequence of $\text{Exp}(1)$ random variables. Set $T_0 = 0$, $X_0 = x$ and given T_n and X_{T_n} inductively define

$$\tau_{n+1} := \inf \left\{ t \geq 0 : \int_0^t \kappa(Y_s^{(n+1)}) ds \geq \xi_{n+1} \right\} \quad (10)$$

where for each $n = 0, 1, 2, \dots$, $(Y_s^{(n+1)})_{s \geq 0}$ is an independent realization of the solution to the SDE (2) started from $Y_0^{(n+1)} = X_{T_n}$. For a careful treatment of defining diffusions on manifolds, the reader is referred to [27].

Then update the path $X_{T_n+s} = Y_s^{(n+1)}$ for $s \in [0, \tau_{n+1})$. Set $T_{n+1} = T_n + \tau_{n+1}$ and independently draw $X_{T_{n+1}} \sim \mu_{T_{n+1}}$, which only depends on the path before time T_{n+1} .

Remark 2.3. The practical simulation of this process, namely of the SDE dynamics (2) and the killing times (10), may seem difficult at first glance. For the killing times, in the present compact setting simulation is actually straightforward, since one can directly employ Poisson thinning (see [13, Chater VI.2.4]) as κ is bounded. Even in noncompact spaces simulation of such times may be performed without error through *layered processes*; this is the case for the ScaLE algorithm [23]. To simulate the SDE, it turns out that through the techniques of *exact simulation* [6], in many settings the SDE (2) can be simulated on fixed time horizons without error and without resorting to discretization. A thorough computational analysis of the resulting ReScaLE algorithm, including selection of the weights η (c.f. Remark 2.2), applications to large data sets and comparisons to existing methods will be the subject of a forthcoming piece of work.

2.2. Fixed rebirth processes

We now define the fixed rebirth processes and derive some useful properties. These will be crucial later for defining the deterministic semiflow. It will be convenient to work on M with the measure

$$\Gamma(dy) = \gamma(y)dy,$$

where

$$\gamma(y) = \exp(2A(y)).$$

Let us write L^κ for minus the generator of the diffusion Y from (2) killed at rate κ and let $\mathcal{D}(L^\kappa) \subset C(M)$ be its domain. Such f are twice continuously differentiable, and we have

$$-L^\kappa f = \frac{1}{2} \Delta f + \nabla A \cdot \nabla f - \kappa f.$$

Then we have the identity

$$e^{-tL^\kappa} f(x) = \mathbb{E}_x[f(Y_t)1_{\{\tau_\partial > t\}}] = \mathbb{E}_x \left[f(Y_t) e^{-\int_0^t \kappa(Y_s) ds} \right]. \quad (11)$$

Here (and throughout) τ_∂ denotes a general killing time, defined analogously to (10) and Y evolves according to our SDE (2) without any killing. The exponentiation of the operator $-L^\kappa$ is rigorously justified through the spectral theorem for self-adjoint operators, and the resulting (sub-Markovian) semigroup $(e^{-tL^\kappa})_{t \geq 0}$ is strongly continuous on $C(M)$. Details may be found in [12].

Similarly to [19,28], we show the existence of a continuous, positive transition subdensity for the SDE (2) killed at rate κ .

Lemma 2.4. *The SDE (2) killed at rate κ has a C^∞ positive transition subdensity $p^\kappa(t, x, y)$ with respect to Γ , that is,*

$$e^{-tL^\kappa} f(x) = \int_M f(y) p^\kappa(t, x, y) \Gamma(dy).$$

Proof. First, we note that the *unkilled* diffusion has a smooth positive transition density $p^0(t, x, y)$ with respect to Γ . The existence of this density is described briefly in Example 9 of [12, Chapter 1.C], and the details can be found in [7, Chapter II]. In particular, the assumptions of [7] are that the manifold M is a C^∞ compact connected finite-dimensional Riemannian manifold (as in our Assumption 1) and that the drift $-\text{our } \nabla A -$ is a C^∞ vector field, which we are assuming (Assumption 2).

In order to obtain the transition subdensity for the killed diffusion, we make use of (11). By conditioning on the end point we can write

$$e^{-tL^\kappa} f(x) = \mathbb{E}_x [f(Y_t)g(t, x, Y_t)] = \int_M f(y)g(t, x, y)p^0(t, x, y)\Gamma(dy),$$

where for each $t > 0$, $x, y \in M$,

$$g(t, x, y) := \mathbb{E}_x \left[e^{-\int_0^t \kappa(Y_s) ds} \mid Y_t = y \right].$$

By [7, Chapter II.d], since κ is bounded C^∞ , $g(t, x, y)$ is continuous in x, y and in fact since κ is smooth and nonnegative, $g(t, x, y)$ is jointly continuous over t, x, y and smooth as a function of x or y . Since κ is nonnegative and bounded above, we have the bounds $e^{-t\bar{\kappa}} \leq g(t, x, y) \leq 1$. In particular $g(t, x, y)$ is positive, and hence setting for each $t > 0$, $x, y \in M$,

$$p^\kappa(t, x, y) = p^0(t, x, y)g(t, x, y), \quad (12)$$

we obtain the positive C^∞ transition subdensity $p^\kappa(t, x, y)$ with respect to Γ of diffusion (2) killed at rate κ . \square

Lemma 2.5. *The resolvent operator $\mathcal{R} : C(M) \rightarrow C(M)$ given by*

$$\mathcal{R}f(x) = \int_0^\infty dt \int_M \Gamma(dy) p^\kappa(t, x, y)f(y)$$

is a well-defined bounded, positive linear operator.

Remark 2.6. We have defined \mathcal{R} as an operator on the Banach space $C(M)$. Its dual operator acts on the space of finite signed Borel measures on M (cf. [15, Section 7.4]). Following the standard probabilistic notation we will denote its dual action on a measure μ by simply $\mu\mathcal{R}$. That is, $\mu\mathcal{R}$ is the measure defined by

$$\mu\mathcal{R}(f) = \int_M \mu(dx) \mathcal{R}f(x).$$

Proof. \mathcal{R} is clearly linear and maps nonnegative functions to nonnegative functions. It maps continuous functions to continuous functions since $p^\kappa(t, x, y)$ and γ are continuous. Thus \mathcal{R}

is a positive linear operator mapping $C(M) \rightarrow C(M)$. For the constant function $1 : x \mapsto 1$ we have

$$\mathcal{R}1(x) = \int_0^\infty dt \int_M \Gamma(dy) p^\kappa(t, x, y) = \int_0^\infty dt \mathbb{P}_x(\tau_\partial > t) = \mathbb{E}_x[\tau_\partial]. \quad (13)$$

Since we are assuming that the killing rate is everywhere bounded below by $\underline{\kappa}$, it follows that we have the uniform bound over $x \in M$,

$$\mathbb{E}_x[\tau_\partial] \leq 1/\underline{\kappa}.$$

Thus since \mathcal{R} is positive, it follows that \mathcal{R} is bounded. \square

We note for future reference that

$$\frac{1}{\bar{\kappa}} \leq \inf_{x \in M} \mathcal{R}1(x) \leq \sup_{x \in M} \mathcal{R}1(x) \leq \frac{1}{\underline{\kappa}}. \quad (14)$$

Remark 2.7. Heuristically, the resolvent describes the average cumulative occupation measure of the killed diffusion over a single lifetime.

Fix a probability measure $\mu \in \mathcal{P}(M)$. We now define the fixed rebirth process with rebirth distribution μ , abbreviated to $\text{FR}(\mu)$, to be a Markov process with càdlàg paths, evolving according to the SDE (2) between regeneration events, which occur at rate $\kappa(X_t^\mu)$. At such an event the location is drawn independently from distribution μ . It can be constructed explicitly as in the construction at the end of Section 2.1 in the simpler case when $X_{T_{n+1}} \sim \mu$ for each n . Let $(P_t^\mu)_{t \geq 0}$ denote the semigroup of this process.

Since κ is continuous and bounded and A is smooth, it can be easily shown that $(P_t^\mu)_{t \geq 0}$ is a strongly continuous semigroup on $C(M)$, so the $\text{FR}(\mu)$ is a Feller–Markov process, and so by the Hille–Yosida theorem (Theorem 1.7, [18]) it has an infinitesimal generator $-L_\mu$ defined on a dense domain $\mathcal{D}(L_\mu) \subset C(M)$.

The action of the generator $-L_\mu$ on smooth functions agrees with (9). Since κ is bounded and continuous, $\mathcal{D}(L_\mu)$ will consist of twice continuously differentiable functions on M , and in fact $\mathcal{D}(L_\mu)$ is independent of μ .

Recall that Y denotes the unkilld process that evolves according to the SDE (2). Since the $\text{FR}(\mu)$ process exhibits a natural renewal behavior, by conditioning on the first arrival time τ_∂ , we see that

$$P_t^\mu f(x) = \int_0^t \mathbb{E}_x \left[\kappa(Y_s) e^{-\int_0^s \kappa(Y_u) du} \right] \mu P_{t-s}^\mu f ds + \mathbb{E}_x \left[f(Y_t) 1_{\{\tau_\partial > t\}} \right]. \quad (15)$$

The second term can be expressed equivalently as in (11).

Lemma 2.8. Given $\mu \in \mathcal{P}(M)$, an invariant measure for the $\text{FR}(\mu)$ process is given by

$$\Pi(\mu)(f) = \frac{\mu \mathcal{R}f}{\mu \mathcal{R}1}. \quad (16)$$

Proof. Let $f \in \mathcal{D}(L_\mu)$, then it follows that $f \in \mathcal{D}(L^\kappa)$. We wish to show that $\mu \mathcal{R}L_\mu f = 0$ (Proposition 9.2 of [17, Chapter 9]). Note by (11) that we can write

$$\mathcal{R}f = \int_0^\infty e^{-tL^\kappa} f dt.$$

Then

$$\begin{aligned} -\mu \mathcal{R} L_\mu f &= -\mu \int_0^\infty e^{-tL^\kappa} L_\mu f \, dt \\ &= \mu \int_0^\infty e^{-tL^\kappa} \left(-L^\kappa f + \kappa \int f(y) \mu(dy) \right) \\ &= \mu \left([e^{-tL^\kappa} f]_0^\infty + \mu(f) \int_0^\infty e^{-tL^\kappa} \kappa \, dt \right) \end{aligned}$$

where we used the backward equation (see, for instance, [17, Chapter 1], Proposition 1.5),

$$\frac{d}{dt}(e^{-tL^\kappa} f) = -e^{-tL^\kappa} L^\kappa f.$$

Note that by Tonelli's theorem we can exchange the order of integration to find

$$\int_0^\infty e^{-tL^\kappa} \kappa \, dt = \mathbb{E} \left[\int_0^{\tau_\partial} \kappa(Y_s) \, ds \right] = \mathbb{E}[\xi] = 1$$

where $\xi \sim \text{Exp}(1)$ by the definition of our killing construction.

Thus putting the terms together we have that

$$-\mu \mathcal{R} L_\mu f = \mu(-f + \mu(f)1) = -\mu(f) + \mu(f) = 0.$$

Thus it follows that $\Pi(\mu)$, which is the normalized version of the measure $\mu \mathcal{R}$, is an invariant probability measure for the $\text{FR}(\mu)$ process. \square

Proposition 2.9. *We have the bound*

$$\|v_1 P_t^\mu - v_2 P_t^\mu\|_1 \leq 2e^{-t\kappa}$$

for any $v_1, v_2, \mu \in \mathcal{P}(M)$. In particular choosing $v_2 = \Pi(\mu)$ gives the bound

$$\|v P_t^\mu - \Pi(\mu)\|_1 \leq 2e^{-t\kappa} \quad (17)$$

for any $v, \mu \in \mathcal{P}(M)$. It follows that $\Pi(\mu)$ is the unique invariant probability measure for the $\text{FR}(\mu)$ process.

Proof. This will follow straightforwardly from the coupling inequality, see, for instance, [25, Section 4.1]. This states that $\|\mathcal{L}(X) - \mathcal{L}(Y)\|_1 \leq 2\mathbb{P}(X \neq Y)$ for a coupling \mathbb{P} of random variables X, Y with laws $\mathcal{L}(X), \mathcal{L}(Y)$ respectively.

Since we are assuming the killing rate κ is bounded below by $\underline{\kappa}$ and the rebirth distribution is fixed, we can couple two processes started from different initial distributions at the first arrival time of a homogeneous Poisson process with rate $\underline{\kappa}$. \square

Lemma 2.10. *The map $\Pi : \mathcal{P}(M) \rightarrow \mathcal{P}(M)$ is continuous in the topology of weak-* convergence.*

Proof. We know that \mathcal{R} is a bounded linear operator on $C(M)$ by Lemma 2.5. Thus as noted in Remark 2.6 it acts dually on the space of finite signed Borel measures, and is continuous on the dual space. So it must also be weak-* continuous. Continuity of $\mu \mapsto \Pi(\mu)$ follows. \square

Remark 2.11. In fact, since $\mathcal{P}(M)$ is a compact metric space it follows that $\Pi : \mathcal{P}(M) \rightarrow \mathcal{P}(M)$ is uniformly continuous.

Proposition 2.12. $\mu \in \mathcal{P}(M)$ satisfies the fixed point equation

$$\mu = \Pi(\mu) \quad (18)$$

if and only if μ is quasi-stationary for the diffusion Y killed at rate κ .

There exists a unique quasi-stationary distribution π for the diffusion Y killed at rate κ . Furthermore, π has a strictly positive C^∞ density with respect to the Riemannian measure, which will also be denoted by π .

Proof. Suppose μ is invariant for L_μ . This means that for all smooth f ,

$$\mu L_\mu f = 0;$$

that is, by (16)

$$\frac{1}{2} \mu(\Delta f) + \mu(\nabla A \cdot \nabla f) + \mu(\kappa)\mu(f) - \mu(\kappa f) = 0,$$

which is equivalent to

$$\mu L^\kappa f = \mu(\kappa)\mu(f).$$

Since $\mu(\kappa) > 0$ this tells us that μ is a quasi-stationary distribution for X .

Conversely, suppose μ is quasi-stationary for Y . Then

$$\mu L^\kappa f = \lambda_0 \mu(f)$$

for all smooth f and some $\lambda_0 > 0$. Then choosing $f \equiv 1$, we find that $\mu(\kappa) = \lambda_0$, from which it follows that $\mu L_\mu f = 0$. Hence μ is stationary for L_μ .

Existence and uniqueness of the quasi-stationary distribution π follows from Theorem 1.1 of [9]. For each $t > 0$, the transition subdensity $p^\kappa(t, x, y)$ is strictly bounded away from 0; we have that for fixed $t > 0$, we can find some $\epsilon > 0$ such that $p^\kappa(t, x, y) > \epsilon$, for all $x, y \in M$. Then the Assumption (A) of Theorem 1.1 of [9] is easily verified.

In addition, it follows from the existence of the transition subdensity $p^\kappa(t, x, y)$ that π must also be absolutely continuous with respect to the Riemannian measure: Since it is the quasi-stationary distribution of the diffusion Y , by basic properties of quasi-stationary distributions (e.g. Theorem 2.2 of [11]) there exists some $\lambda_0^\kappa > 0$ such that for each measurable $A \subset M$ and $t > 0$,

$$\pi(A) = \mathbb{P}_\pi(Y_t \in A \mid \tau_\partial > t) = \int_M \pi(dx) \int_A dy \gamma(y) e^{t\lambda_0^\kappa} p^\kappa(t, x, y).$$

In particular this implies that π is absolutely continuous with respect to the Riemannian measure; hence π has a density with respect to it, which we will also denote by π . Thus the density π satisfies for each $t > 0$,

$$\pi(y) = \int_M \pi(dx) e^{t\lambda_0^\kappa} p^\kappa(t, x, y) \gamma(y), \quad (19)$$

for almost every $y \in M$. But since $p^\kappa(t, x, y)$ is positive and smooth, it follows that the density π is continuous – so (19) holds for all $y \in M$ – and then the density π is smooth, and also positive everywhere. \square

Remark 2.13. As noted in Section 2, in the context of Monte Carlo the quasi-stationary distribution π will be chosen to coincide with a distribution of interest by choosing κ according to (7).

3. Deterministic flow

3.1. Basic properties

We are now in a position to define the deterministic measure-valued flow that will characterize the asymptotic behavior of the normalized occupation measures $(\mu_t)_{t \geq 0}$.

Recall, as in [2], that on a metric space E a *semiflow* Φ is a jointly continuous map

$$\begin{aligned}\Phi : \mathbb{R}_+ \times E &\rightarrow E, \\ (t, x) &\mapsto \Phi(t, x) = \Phi_t(x)\end{aligned}$$

such that Φ_0 is the identity on E and $\Phi_{t+s} = \Phi_t \circ \Phi_s$ for all $s, t \in \mathbb{R}_+$.

We would like to define a semiflow Φ on the space $E = \mathcal{P}(M)$ of probability measures with the topology of weak-* convergence, which is a metric space. In particular we want $t \mapsto \Phi_t(\mu)$ to solve the measure-valued ordinary differential equation (ODE)

$$\dot{v}_t = -v_t + \Pi(v_t), \quad v_0 = \mu, \quad (20)$$

in the weak sense, meaning that for any test function $f \in C(M)$

$$\Phi_t(\mu)f = \mu f + \int_0^t (-\Phi_s(\mu)f + \Pi(\Phi_s(\mu))f) ds.$$

We define such a semiflow by adapting the approach of [4, Section 5] to our present setting. As noted in Lemma 2.5, the operator \mathcal{R} is a bounded linear operator mapping from the Banach space $C(M)$ to itself. This allows us to define, for any $t \geq 0$, the bounded linear operator $e^{t\mathcal{R}} : C(M) \rightarrow C(M)$, whose dual acts on the space of finite signed Borel measures, equipped with the total-variation norm.

This allows us to define, for each $t \geq 0$, the probability measures

$$\tilde{\Phi}_t(\mu) := \frac{\mu e^{t\mathcal{R}}}{\mu e^{t\mathcal{R}} 1}.$$

The map $t \mapsto \tilde{\Phi}_t(\mu)$ satisfies the weak measure-valued ODE

$$\dot{v}_t = v_t \mathcal{R} - (v_t \mathcal{R} 1) v_t = (v_t \mathcal{R} 1)(-v_t + \Pi(v_t)), \quad v_0 = \mu.$$

To get a solution to our actual ODE (20) we employ a suitable time change, imitating [8] and [4]. Similarly to [4], for $t \geq 0$ set

$$s_\mu(t) := \int_0^t \tilde{\Phi}_s(\mu) \mathcal{R} 1 ds$$

so $\dot{s}_\mu(t) = \tilde{\Phi}_t(\mu) \mathcal{R} 1 = \mathbb{E}_{\tilde{\Phi}_t(\mu)}[\tau_\partial] > 0$, so $t \mapsto s_\mu(t)$ is strictly increasing.

Lemma 3.1. $\dot{s}_\mu(t) \geq 1/\bar{\kappa} > 0$ for all $\mu \in \mathcal{P}(M)$ and $t > 0$. Thus in particular $s_\mu(t) \rightarrow \infty$ as $t \rightarrow \infty$ for any $\mu \in \mathcal{P}(M)$.

Proof. Immediate since the function $x \mapsto \mathbb{E}_x[\tau_\partial]$ is uniformly bounded below by $1/\bar{\kappa}$ since the killing rate is bounded above by $\bar{\kappa}$. \square

For a fixed $\mu \in \mathcal{P}(M)$, since s_μ is a strictly increasing, differentiable map $\mathbb{R}_+ \rightarrow \mathbb{R}_+$ we can define an inverse mapping τ_μ , and compose

$$\Phi_t(\mu) = \tilde{\Phi}_{\tau_\mu(t)}(\mu).$$

Recall that we have equipped the space $\mathcal{P}(M)$ with the weak-* topology, which is a compact metric space.

Proposition 3.2. *Φ is an injective semiflow on $\mathcal{P}(M)$, and for each $\mu \in \mathcal{P}(M)$ $t \mapsto \Phi_t(\mu)$ is the unique weak solution to (20).*

Proof. This is identical to the proof of [4, Proposition 5.1]. \square

3.2. Stability of π

Recall that we are interested in working with respect to the measure $\Gamma(dy) = \gamma(y)dy$ where $\gamma = \exp(2A)$. From Proposition 2.12, we see that there is a (unique) quasi-stationary distribution π which is a fixed point of Π , which also has a density with respect to the Riemannian measure, which we also denote by π . Set

$$\varphi := \pi/\gamma$$

which is a smooth function.

We now list some basic facts about π and φ which can all be verified through routine manipulations. There exists $\lambda_0^\kappa > 0$, which describes the asymptotic rate of killing, such that

$$L^\kappa \varphi = \lambda_0^\kappa \varphi.$$

φ is bounded above and bounded away from 0 since by Proposition 2.12, π is both bounded above and bounded away from zero.

Writing

$$\beta := \frac{1}{\lambda_0^\kappa}$$

it follows that

$$\begin{aligned} \mathcal{R}\varphi &= \beta\varphi, \\ \pi\mathcal{R} &= \beta\pi \end{aligned} \tag{21}$$

where this final identity holds both in terms of π as a measure, and pointwise as density functions.

We wish now to analyze the asymptotic behavior of the semiflow Φ defined in Section 3. To do this we will derive a drift condition for a reweighted version of the Markov process.

Since φ is bounded above and away from zero, consider the bounded linear operator $L : C(M) \rightarrow C(M)$ given by

$$L = \frac{1}{\varphi}(\mathcal{R} - \beta)\varphi.$$

Thus by exponentiation L generates a Markov semigroup $(\mathcal{K}^t)_{t \geq 0}$. For each $t \geq 0$,

$$\mathcal{K}^t f = \frac{1}{\varphi} \exp(t(\mathcal{R} - \beta))(\varphi f).$$

We can also define the kernels $(\tilde{\mathcal{K}}^t)_{t \geq 0}$,

$$\tilde{\mathcal{K}}^t := \varphi \mathcal{K}^t(f/\varphi) = \exp(t(\mathcal{R} - \beta))f.$$

For the process defined by $(\mathcal{K}^t)_{t \geq 0}$ it can easily be seen from (21) that the measure $\pi\varphi$ given by $(\pi\varphi)(f) = \int f(x)\varphi(x)\pi(x)dx$ is an invariant measure. Since φ is bounded, we see that $\pi\varphi$ is in fact a finite measure. In what proceeds, without loss of generality we rescale φ so that $\pi\varphi(1) = 1$.

We would like to show that this process is V -uniformly ergodic, with $V = 1/\varphi$. We do this using a drift condition from Theorem 5.2 of [14]: for a continuous-time irreducible aperiodic Markov process with extended generator L , assume it satisfies for constants $b, c > 0$ and petite set C

$$LV \leq -cV + b1_C, \quad (22)$$

then the process is V -uniformly ergodic. Heuristically, a petite set is a set from which the Markov process leaves with a common minorizing measure. The precise definition is carefully presented in [14, Section 3]. For our present purposes, since we have a positive continuous transition subdensity $p^\kappa(t, x, y)$, it is enough to note that compact sets (and hence the entire space) are petite.

Proposition 3.3. *The drift condition (22) holds.*

Proof. Recall we have set $V = 1/\varphi$. Then for $\beta = 1/\lambda_0^\kappa$,

$$LV = \frac{1}{\varphi}(\mathcal{R} - \beta)1 = \frac{1}{\varphi}(\mathbb{E}[\tau_\partial] - \beta) = -\beta V + V\mathbb{E}[\tau_\partial] \leq -\beta V + b$$

where b is an upper bound for $V(x)\mathbb{E}_x[\tau_\partial]$ for all x , which exists since V is bounded above and κ is bounded away from 0. Note our entire space is compact, hence petite. Thus the drift condition holds. \square

Remark 3.4. We note that since the entire space M is compact hence petite, the drift condition (22) can be trivially satisfied by choosing $V \equiv 1$, the constants $b = c = 1$ and $C = M$. However we have kept the choice $V = \varphi$ in the proof since this choice is suggestive of how one might generalize this to noncompact spaces.

By [14, Theorem 5.2] this implies V -uniform ergodicity: There exist constants D and $0 \leq \rho < 1$ such that for all $x \in M$:

$$\sup_{|g| \leq V} |\mathcal{K}^t(x, g) - \pi\varphi(g)| \leq V(x)D\rho^t.$$

Multiplying through by $\varphi(x)$ and relabeling φg as f , we see that the condition $|g| \leq V = 1/\varphi$ is equivalent to $|f| \leq 1$, hence proving uniform ergodicity: For any $x \in M$

$$\sup_{|f| \leq 1} |\tilde{\mathcal{K}}^t(x, f) - \varphi(x)\pi(f)| \leq D\rho^t.$$

Thus we will have, for any initial distribution μ ,

$$\sup_{|f| \leq 1} |\mu\tilde{\mathcal{K}}^t f - \mu(\varphi)\pi(f)| \leq D\rho^t. \quad (23)$$

Proposition 3.5. *We have convergence $\tilde{\Phi}_t(\mu) \rightarrow \pi$ as $t \rightarrow \infty$ in total variation distance, uniformly in μ .*

Proof. Let $\varphi_* := \min_{x \in M} \varphi(x)$, which is positive, since φ is a continuous positive function on a compact set. We find that for any $t \geq (\log 2D - \log \varphi_*) / \log \rho^{-1}$ we have by (23)

$$\mu \tilde{\mathcal{K}}^t 1 \geq \mu(\varphi) - D\rho^t \geq \frac{\varphi_*}{2}.$$

Since

$$\tilde{\Phi}_t(\mu) = \frac{\mu \tilde{\mathcal{K}}^t}{\mu \tilde{\mathcal{K}}^t 1},$$

we have then for any probability measure μ and continuous f with $|f| \leq 1$

$$\begin{aligned} |\tilde{\Phi}_t(\mu)f - \pi(f)| &\leq \left(\mu \tilde{\mathcal{K}}^t 1\right)^{-1} \left[|\mu \tilde{\mathcal{K}}^t f - \mu(\varphi)\pi(f)| + |\pi(f)| \cdot |\mu(\varphi) - \mu \tilde{\mathcal{K}}^t 1| \right] \\ &\leq \frac{2}{\varphi_*} \cdot [1 + |\pi(f)|] D\rho^t \\ &\leq \frac{4D}{\varphi_*} \rho^t. \quad \square \end{aligned}$$

Finally, we see that this convergence carries over to the semiflow Φ .

Theorem 3.6. *We have convergence $\Phi_t(\mu) \rightarrow \pi$ as $t \rightarrow \infty$ uniformly in μ in total variation distance.*

Proof. This follows from Proposition 3.5 and the fact that $\dot{s}_\mu(t)$ is bounded above by $1/\kappa$ uniformly in μ and t , as in Lemma 2.5. The boundedness of this derivative ensures that its inverse τ_μ satisfies for all $\mu \in \mathcal{P}(M)$ and $t \geq 0$

$$\tau_\mu(t) \geq \kappa t. \quad \square$$

Remark 3.7. In the language of [2], this shows that π is a *global attractor* of the semiflow Φ .

4. Asymptotic pseudo-trajectories

4.1. Basic properties and definitions

We now wish to relate the stochastic behavior of the normalized weighted empirical measures

$$\mu_t = \frac{r\mu_0 + \int_0^t \eta_s \delta_{X_s} ds}{r + \int_0^t \eta_s ds}$$

to the deterministic behavior of the flow defined in Section 3.

Definition. ([2, Section 3]) For a metric space (E, d) , given a semiflow Φ on E , a continuous function $w : [0, \infty) \rightarrow E$ is an *asymptotic pseudo-trajectory* of Φ if for all $T > 0$,

$$\lim_{t \rightarrow \infty} \sup_{s \in [0, T]} d(w(t+s), \Phi_s(w(t))) = 0.$$

Recall that by Assumption 4, since g is continuously differentiable, strictly increasing and $g(t) \rightarrow \infty$ as $t \rightarrow \infty$, it is a diffeomorphism of $\mathbb{R}_+ \rightarrow \mathbb{R}_+$. Thus we can define its inverse

function $g^{-1} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is also continuously differentiable and satisfies $g^{-1}(t) \rightarrow \infty$ as $t \rightarrow \infty$. Let

$$\zeta_t := \mu_{h(t)}$$

for all $t \geq 0$, where we defined $h(t) := g^{-1}(re^t - r)$. We will show that almost surely $t \mapsto \zeta_t$ is an asymptotic pseudo-trajectory of the semiflow Φ defined in Section 3.

Since

$$\frac{d}{dt}\mu_t = \alpha_t(-\mu_t + \delta_{X_{t-}}), \quad (24)$$

applying the chain rule and product rule for derivatives yields

$$\frac{d}{dt}\zeta_t = (-\zeta_t + \Pi(\zeta_t)) + (\delta_{X_{h(t)-}} - \Pi(\zeta_t)). \quad (25)$$

Looking at the first bracket, we recognize the flow from Section 3. Thus to formally show that ζ_t approximates the flow, we need to control the second bracket. We also note here for future reference that

$$\begin{aligned} \frac{\partial}{\partial t}\Pi(\mu_t)f &= \alpha_t\left(-\frac{\mu_t\mathcal{R}f}{\mu_t\mathcal{R}1} + \frac{\mathcal{R}f(X_{t-})}{\mu_t\mathcal{R}1} + \frac{\mu_t\mathcal{R}f}{\mu_t\mathcal{R}1} - \frac{\mu_t\mathcal{R}f \cdot \mathcal{R}1(X_{t-})}{(\mu_t\mathcal{R}1)^2}\right) \\ &= \frac{\alpha_t}{\mu_t\mathcal{R}1}\left(\mathcal{R}f(X_{t-}) - \mathcal{R}1(X_{t-})\Pi(\mu_t)f\right). \end{aligned} \quad (26)$$

We formalize this intuition by the approach of [5, Proposition 3.5] and [20, Lemma 5.4]. It is proven in Theorem 3.2 of [2] that asymptotic pseudo-trajectories must be uniformly continuous. Conversely, Theorem 3.2 of [2] also tells us that a uniformly continuous path ζ is an asymptotic pseudo-trajectory if and only if every limit point of the time shifts $\Theta^t\zeta$ in the topology of uniform convergence on compact sets is itself a trajectory of the flow. We define $\mathcal{M}(M)$ to be the space of Borel signed measures on M , equipped with the weak-* topology, which can be metrized analogously to (5). Let $C(\mathbb{R}_+, \mathcal{P}(M))$ and $C(\mathbb{R}_+, \mathcal{M}(M))$ be the spaces of continuous paths mapping \mathbb{R}_+ into $\mathcal{P}(M)$ and $\mathcal{M}(M)$ respectively, each equipped with the topology of uniform convergence on compact subsets of \mathbb{R}_+ . As usual, for each $t \geq 0$ we define $\Theta^t : C(\mathbb{R}_+, \mathcal{P}(M)) \rightarrow C(\mathbb{R}_+, \mathcal{P}(M))$ to be the shift map given by

$$[\Theta^t\zeta]_s = \zeta_{t+s}, \quad s \geq 0.$$

Defining the *retraction* $\hat{\Phi} : C(\mathbb{R}_+, \mathcal{P}(M)) \rightarrow C(\mathbb{R}_+, \mathcal{P}(M))$ as in [2] and [5, Proposition 3.5] by

$$\hat{\Phi}(\zeta)(s) = \Phi_s(\zeta_0), \quad s \geq 0,$$

showing that ζ is an asymptotic pseudo-trajectory of Φ is then equivalent to showing that the limit points of $\{\Theta^t\zeta\}_{t \geq 0}$ are fixed points of the retraction $\hat{\Phi}$.

We also define in analogue to [5] the operator $L_F : C(\mathbb{R}_+, \mathcal{P}(M)) \rightarrow C(\mathbb{R}_+, \mathcal{M}(M))$ by

$$L_F(v)(s) = v(0) + \int_0^s F(v(u))du, \quad s \geq 0,$$

with $F(v) := -v + \Pi(v)$, which will be used in the subsequent proof.

Define the collection of signed measures

$$\epsilon_t(s) := \int_t^{t+s} (\delta_{X_{h(u)-}} - \Pi(\zeta_u))du$$

for all $t, s \geq 0$. We note that for each $t \geq 0$, $\epsilon_t(\cdot) \in C(\mathbb{R}_+, \mathcal{M}(M))$.

Theorem 4.1. Suppose $\zeta : \mathbb{R}_+ \rightarrow \mathcal{P}(M)$ is a continuous path as described above. Then ζ is an asymptotic pseudo-trajectory for Φ if and only if we have the following condition:

For any $T > 0$ and smooth $f \in C(M)$

$$\lim_{t \rightarrow \infty} \sup_{s \in [0, T]} |\epsilon_t(s)f| = 0. \quad (27)$$

Remark 4.2. Theorem 3.2 of [2] assumes relative compactness of the image of the path ζ . By Prokhorov's theorem (see [15, Theorem 11.5.4]), relative compactness in $\mathcal{P}(M)$ is equivalent to tightness, which trivially holds in our present compact setting.

Proof. Given continuous f with $\|f\|_\infty \leq 1$, we have

$$|\zeta_{t+s}f - \zeta_t f| \leq 2|s|$$

since

$$\frac{d(\zeta_t f)}{dt} = -\zeta_t f + f(X_{h(t-)}).$$

Hence we have uniform continuity.

Suppose the condition (27) holds for any $T > 0$ and smooth $f \in C(M)$. This says that $\epsilon_t(\cdot)$ converges to 0 in $C(\mathbb{R}_+, \mathcal{M}(M))$. Consider the family $\{\Theta^t \zeta\}_{t \geq 0}$. Suppose $\tilde{\zeta}$ is a limit point of this family in $C(\mathbb{R}_+, \mathcal{P}(M))$. Analogously to [5], we can use (25) and the definition of ϵ_t to write

$$\Theta^t \zeta = L_F(\Theta^t \zeta) + \epsilon_t(\cdot). \quad (28)$$

Since we are assuming precisely (27), and since L_F is continuous, taking $t \rightarrow \infty$ in (28) along the appropriate subsequence we obtain $\tilde{\zeta} = L_F(\tilde{\zeta})$. This shows that the limit path $\tilde{\zeta}$ is a fixed point of L_F , that is,

$$\tilde{\zeta}_s = \tilde{\zeta}_0 + \int_0^s F(\tilde{\zeta}_u) du, \quad s \geq 0.$$

By uniqueness of the flow this implies that $\tilde{\zeta}_s = \Phi_s(\tilde{\zeta}_0)$ for all $s \geq 0$, that is, $\tilde{\zeta}$ is a fixed point of the retraction $\hat{\Phi}$. This concludes the proof of the sufficiency of condition (27).

Conversely, suppose ζ is an asymptotic pseudo-trajectory for Φ . By definition, this means that for each $T > 0$, as $t \rightarrow \infty$,

$$\sup_{s \in [0, T]} d_w(\zeta_{t+s}, \Phi_s(\zeta_t)) \rightarrow 0. \quad (29)$$

By the representation (28) we would like to show that for each $T > 0$,

$$\sup_{s \in [0, T]} d_w(\Theta^t \zeta(s), L_F(\Theta^t \zeta)(s)) = \sup_{s \in [0, T]} d_w\left(\zeta_{t+s}, \zeta_t + \int_0^s F(\zeta_{t+u}) du\right) \rightarrow 0$$

By (29), it is sufficient to control $d_w(L_F(\Theta^t \zeta_s), \Phi_s(\zeta_t))$ uniformly over $s \in [0, T]$.

By the definition of the flow Φ , this will follow if we can control

$$\sup_{s \in [0, T]} d_w\left(\int_0^s F(\zeta_{t+u}) du, \int_0^s F(\Phi_u(\zeta_t)) du\right). \quad (30)$$

Since F is uniformly continuous (since $\Pi : \mathcal{P}(M) \rightarrow \mathcal{P}(M)$ is uniformly continuous; Remark 2.11), (30) can be made arbitrarily small by choosing t sufficiently large because of (29).

Thus by (28) we conclude that $\epsilon_t \rightarrow 0$ in $C(\mathbb{R}_+, \mathcal{M}(M))$, that is, (27) must hold for any $T > 0$ and smooth $f \in C(M)$ as required. \square

The goal now is to establish the requirements of Theorem 4.1 almost surely. So we need to establish (27).

4.2. Poisson equation

To show that (27) holds, inspired by [5, Section 5.2] and [4, Section 6.2], we will measure the discrepancy between the cumulative occupancy and the quasi-stationary distribution by a solution to the Poisson equation

$$-f + \Pi(\mu)f = -L_\mu g. \quad (31)$$

Fix any $\mu \in \mathcal{P}(M)$. Define for $f \in C(M)$,

$$Q_\mu f := \int_0^\infty (P_t^\mu f - \Pi(\mu)f) dt.$$

Lemma 4.3. For any $\mu \in \mathcal{P}(M)$, Q_μ is a bounded linear operator mapping $C(M)$ to itself, and

$$\|Q_\mu f\|_\infty \leq \frac{2\|f\|_\infty}{\kappa}$$

for any $\mu \in \mathcal{P}(M)$ and $f \in C(M)$.

Proof. $Q_\mu f$ is continuous since f is continuous and (P_t^μ) is Feller. From (17), we have that

$$|Q_\mu f(x)| \leq \int_0^\infty \|P_t^\mu(x) - \Pi(\mu)\|_1 \|f\|_\infty dt \leq 2\|f\|_\infty \int_0^\infty e^{-t\kappa} dt = \frac{2\|f\|_\infty}{\kappa}$$

uniformly over $x \in M$ and $\mu \in \mathcal{P}(M)$. \square

Recall that $-L_\mu$ is the generator of the FR(μ) process. By the basic properties of infinitesimal generators (see, for instance, [18]), we have that for any $t > 0$, $f \in C(M)$, that $\int_0^t P_s^\mu f ds \in \mathcal{D}(L_\mu)$, and $(-L_\mu) \int_0^t P_s^\mu f ds = P_t^\mu f - f$. Then

$$P_t^\mu f = f + (-L_\mu) \int_0^t P_s^\mu f ds. \quad (32)$$

Proposition 4.4. Given $f \in C(M)$, $Q_\mu f \in \mathcal{D}(L_\mu)$, and $g = Q_\mu f$ solves the Poisson equation (31).

Proof. Without loss of generality, assume $\Pi(\mu)f = 0$; just replace f with $f - \Pi(\mu)f$. Then taking $t \rightarrow \infty$ in (32) and using Proposition 2.9 we see by closure of $-L_\mu$ that $Q_\mu \in \mathcal{D}(L_\mu)$, and that $\Pi(\mu)f = f - L_\mu Q_\mu f$. Thus (31) is satisfied by $g = Q_\mu f$. \square

Remark 4.5. By basic properties of generators of Feller processes (cf. [24, Section VII.1]), we similarly have that for $f \in \mathcal{D}(L_\mu)$,

$$-f + \Pi(\mu)f = -Q_\mu L_\mu f. \quad (33)$$

Remark 4.6. Note that we can think of the operator

$$K_\mu : f \mapsto f - \Pi(\mu)f$$

as a projection operator since it is linear and idempotent ($K_\mu^2 = K_\mu$).

4.3. Bounding the discrepancy

Using our solution to the Poisson equation, we now rewrite the discrepancy term, and decompose it using Itô's formula. Making the change of variables $u \leftarrow h(u)$ we write

$$\begin{aligned}\epsilon_t(s)f &= \int_t^{t+s} (f(X_{h(u-)})) - \Pi(\zeta_u)f \, du \\ &= \int_{h(t)}^{h(t+s)} (-L_{\mu_u})Q_{\mu_u}f(X_{u-})\alpha_u \, du.\end{aligned}$$

Consider now a time-dependent function $f : \mathbb{R} \times M \rightarrow \mathbb{R}$, written as $f_s(x)$. For simplicity we will take the notation ∇ and ∇^2 to refer to the gradient with respect to the coordinates of M , and f'_s to be the time derivative $\partial f_t / \partial t \big|_{t=s}$. We now apply Itô's formula for general semimartingales (as formulated, for instance, as Theorem 14.2.4 of [10]), taking advantage of the fact that all quadratic covariation terms are 0 for the Brownian motion that is driving our process X_t :

$$\begin{aligned}f_t(X_t) - f_0(X_0) &= \int_0^t \nabla f_u(X_{u-}) \cdot dX_u + \frac{1}{2} \int_0^t \nabla^2 f_u(X_{u-}) du + \int_0^t f'_u(X_{u-}) du \\ &\quad + \sum_{0 < u \leq t} (f_u(X_u) - f_u(X_{u-}) - \nabla f_u(X_{u-}) \cdot \Delta X_u).\end{aligned}$$

Using the formula (9) for L_μ we find

$$\begin{aligned}f_t(X_t) - f_0(X_0) &= \int_0^t (-L_{\mu_u})f_u(X_{u-})du + \int_0^t \nabla f_u(X_{u-}) \cdot dW_u + \int_0^t f'_u(X_{u-})du \\ &\quad + \sum_{0 < u \leq t} (f_u(X_u) - f_u(X_{u-})) \\ &\quad - \int_0^t \kappa(X_{u-}) \int (f_u(y) - f_u(X_{u-}))\mu_u(dy)du,\end{aligned}$$

where W_u is a Wiener process on M .

We apply this formula now to the function $f_s(x) = Q_{\mu_s}f(x)\alpha_s$. By Proposition 4.4, $Q_{\mu_s}f \in \mathcal{D}(L_{\mu_s})$, and so in particular $Q_{\mu_s}f$ is twice continuously differentiable. Thus $f_s(x)$ is indeed twice continuously differentiable with respect to x . We rearrange the terms to obtain

$$\begin{aligned}\epsilon_s(t)f &= \int_{h(t)}^{h(t+s)} (-L_{\mu_u})Q_{\mu_u}f(X_{u-})\alpha_u \, du \\ &= \epsilon_s^{(1)}(t)f + \epsilon_s^{(2)}(t)f + \epsilon_s^{(3)}(t)f + \epsilon_s^{(4)}(t)f + \epsilon_s^{(5)}(t)f,\end{aligned}\tag{34}$$

where

$$\begin{aligned}\epsilon_t^{(1)}(s)f &= Q_{\mu_{h(t+s)}}f(X_{h(t+s)-})\alpha_{h(t+s)} - Q_{\mu_{h(t)}}f(X_{h(t)-})\alpha_{h(t)}, \\ \epsilon_t^{(2)}(s)f &= \int_{h(t)}^{h(t+s)} Q_{\mu_u}f(X_{u-})\frac{d\alpha_u}{du}du,\end{aligned}$$

$$\begin{aligned}\epsilon_t^{(3)}(s)f &= - \int_{h(t)}^{h(t+s)} \frac{\partial}{\partial u} \left(\mathcal{Q}_{\mu_u} f \right) (X_{u-}) \alpha_u du, \\ \epsilon_t^{(4)}(s)f &= -N_{h(t+s)}^f + N_{h(t)}^f, \\ \epsilon_t^{(5)}(s)f &= -J_{h(t+s)}^f + J_{h(t)}^f,\end{aligned}$$

and N and J are the local martingales

$$\begin{aligned}N_t^f &:= \int_0^t \nabla \mathcal{Q}_{\mu_u} f(X_{u-}) \alpha_u \cdot dW_u, \\ J_t^f &:= \sum_{0 < u \leq t} \left(\mathcal{Q}_{\mu_u} f(X_u) \alpha_u - \mathcal{Q}_{\mu_u} f(X_{u-}) \alpha_u \right) \\ &\quad - \int_0^t du \kappa(X_{u-}) \int \mu_u(dy) \left(\mathcal{Q}_{\mu_u} f(y) \alpha_u - \mathcal{Q}_{\mu_u} f(X_{u-}) \alpha_u \right).\end{aligned}$$

Theorem 4.7. *The conditions of Theorem 4.1 hold almost surely. This implies that $t \mapsto \zeta_t$ is almost surely an asymptotic pseudo-trajectory for Φ .*

Proof. To establish (27) we will use the decomposition (34) and consider the five error terms individually.

4.3.1. $\epsilon_t^{(1)}(s)f$ and $\epsilon_t^{(2)}(s)f$

Using the bound from Lemma 4.3, we see that

$$\begin{aligned}|\epsilon_t^{(1)}(s)f| &\leq \|\mathcal{Q}_{\mu_{h(t+s)}} f\|_\infty \alpha_{h(t+s)} + \|\mathcal{Q}_{\mu_{h(t)}} f\|_\infty \alpha_{h(t)} \\ &\leq 4\kappa^{-1} \|f\|_\infty (\alpha_{h(t+s)} - \alpha_{h(t)}), \\ |\epsilon_t^{(2)}(s)f| &\leq \int_{h(t)}^{h(t+s)} \|\mathcal{Q}_{\mu_u} f\|_\infty \frac{d\alpha_u}{du} du \\ &= 2\kappa^{-1} \|f\|_\infty (\alpha_{h(t+s)} - \alpha_{h(t)}).\end{aligned}$$

Since $h(t) \rightarrow \infty$ as $t \rightarrow \infty$ and $\alpha_t \rightarrow 0$ as $t \rightarrow \infty$ by Assumption 4, we see that both terms decay to 0 as $t \rightarrow \infty$.

4.3.2. $\epsilon_t^{(3)}(s)f$

By the same argument as in the proof of Lemma 5.5 of [5] we have

$$\frac{\partial}{\partial t} \mathcal{Q}_{\mu_t} = - \left(\frac{\partial}{\partial t} K_{\mu_t} + \mathcal{Q}_{\mu_t} \frac{\partial}{\partial t} (-L_{\mu_t}) \right) \mathcal{Q}_{\mu_t}, \quad (35)$$

where $K_\mu f = f - \Pi(\mu)f$. While our $-L_\mu$ is not the same as their operator A_μ , the same proof holds, since it relies only upon the Poisson equations (31) and (33).

Applying (24) to the definition (9) of L_{μ_t} we obtain

$$-\frac{\partial}{\partial t} L_{\mu_t} f(x) = \kappa(x) \alpha_t \left(f(X_{t-}) - \mu_t(f) \right).$$

From (26) we obtain an upper bound

$$\left\| \alpha_t^{-1} \frac{\partial K_{\mu_t}}{\partial t} \mathcal{Q}_{\mu_t} f \right\|_\infty \leq \left| \frac{1}{\mu_t \mathcal{R}1} \right| \left(\|\mathcal{R} \mathcal{Q}_{\mu_t} f(X_{t-})\|_\infty + |\Pi(\mu_t) \mathcal{R}1(X_{t-})| \right),$$

so by Lemma 4.3 and the bounds (14) we see that there is a constant C (depending on the upper and lower bounds on κ) such that

$$\sup_{x \in M} \left| \frac{\partial Q_{\mu_t}}{\partial t} f(x) \right| \leq C \|f\|_{\infty} \alpha_t.$$

This implies that

$$|\epsilon_t^{(3)}(s)f| \leq C \|f\|_{\infty} \int_{h(t)}^{\infty} \alpha_u^2 du.$$

Since the definite integral from 0 to ∞ is finite by Assumption 4, it follows immediately that $\epsilon_t^{(3)}(s)f$ decays to 0 as $t \rightarrow \infty$.

4.3.3. $\epsilon_t^{(4)}(s)f$

We now turn to the first of the two martingale terms. Similarly to Proposition 5.3 of [20], our goal is to control the quadratic variation and then apply the Burkholder–Davis–Gundy inequality.

The quadratic variation of the martingale N_t^f is bounded by

$$\int_0^t \|\nabla Q_{\mu_s} f(X_s)\|_{\infty}^2 \alpha_u^2 ds.$$

We can bound this by means of the inequality

$$\|\nabla P_t^{\mu} f\|_{\infty} \leq \frac{C_1 \|f\|_{\infty}}{\sqrt{t}}, \quad \text{for } t \text{ sufficiently small,} \quad (36)$$

for some constant C_1 . Such an inequality was shown to hold in the proof of Lemma 5.1 of [5] for diffusive processes on compact manifolds without jumps. To see that this inequality also holds in the present setting, consider the expression for the semigroup (15):

$$P_t^{\mu} f(x) = \int_0^t \mathbb{E}_x [\kappa(Y_s) 1_{\{\tau_0 > s\}}] \mu P_{t-s}^{\mu} f ds + \mathbb{E}_x [f(Y_t) 1_{\{\tau_0 > t\}}].$$

With this representation, it suffices to show that an inequality of the form (36) for the killed semigroup given in Lemma 2.4. This will ensure that our resulting bound (36) is uniform over μ .

By (12), since $|g(t, x, y)| \leq 1$ uniformly, and since we know that such an inequality holds for diffusive processes on compact manifolds without jumps, it suffices to show that

$$\|\nabla_x g(t, \cdot, y)\|_{\infty} \leq \frac{C}{\sqrt{t}} \quad (37)$$

for some constant C uniformly over y , for all t sufficiently small. By the Girsanov–Cameron–Martin formula [16] we can write

$$g(t, x, y) = \mathbb{Q}_t \left[e^{-\int_0^t \kappa(Y_s) ds} G(Y) \right]$$

where for $s \in [0, t]$, $Y_s = (1 - s/t)x + (s/t)y + \hat{W}_s$, where under \mathbb{Q}_t , \hat{W} is a standard Brownian bridge with $\hat{W}_0 = \hat{W}_t = 0$. $G(Y)$ is given by

$$G(Y) = \exp \left(A(y) - A(x) - \frac{1}{2} \int_0^t \Delta A(Y_s) ds - \frac{1}{2} \int_0^t \|\nabla A(Y_s)\|^2 ds \right),$$

since our drift is assumed to be of a gradient form ∇A . Having now expressed Y with explicit dependence upon the initial position x , we can calculate $\nabla_x g(t, \cdot, y)$. Since the functions κ

and A are assumed smooth (hence they and their derivatives are bounded uniformly), we can conclude that a bound of the form (37) holds. In fact, the bound we obtain is actually uniform over small t , and does not blow up as $t \rightarrow 0$.

Armed with (36), we then have

$$\int_0^\infty \|\nabla P_t^\mu f\|_\infty dt \leq C_2 \|f\|_\infty, \quad (38)$$

for some constant C_2 , obtained by considering separately the integrals over $(0, t_0)$ and (t_0, ∞) . Applying (36) bounds the former, while the semigroup property allows the latter piece to be bounded by

$$\begin{aligned} \int_{t_0}^\infty \|\nabla P_t^\mu f\|_\infty dt &\leq \int_0^\infty \|\nabla P_{t_0}^\mu (P_t^\mu (K_\mu f))\|_\infty dt \\ &\leq \frac{C_1}{\sqrt{t_0}} \int_0^\infty \|P_t^\mu (K_\mu f)\|_\infty dt. \end{aligned}$$

Here we replaced f with $K_\mu f = f - \Pi(\mu)f$ since $\Pi(\mu)f$ is a constant so $\nabla \Pi(\mu)f = 0$. This final term can be bounded just as in Lemma 4.3.

Since $\nabla Q_\mu f = \int_0^\infty \nabla P_t^\mu dt$, (38) immediately implies a universal bound

$$\|\nabla Q_{\mu_u} f\|_\infty \leq C_3 \|f\|_\infty.$$

for smooth f . This gives us a bound for the quadratic variation of

$$C_3^2 \|f\|_\infty^2 \int_0^t \alpha_s^2 ds.$$

The Burkholder–Davis–Gundy inequality then implies the existence of a constant C_4 such that for any $\delta > 0$

$$\mathbb{P}_x \left(\sup_{s \in [0, T]} |\epsilon_t^{(4)}(s)f| \geq \delta \right) \leq \frac{C_4 \|f\|_\infty^2}{\delta^2} \int_{h(t)}^\infty \alpha_s^2 ds.$$

Using (8) from Assumption 4, it follows from the Borel–Cantelli lemma that almost surely, for any $\delta > 0$ and $T > 0$,

$$\limsup_{n \rightarrow \infty} \sup_{s \in [0, T]} |\epsilon_n^{(4)}(s)f| \leq \delta. \quad (39)$$

Thus for any $\delta > 0$ we may find N sufficiently large so that for any $t \geq N - T$,

$$\sup_{s \in [0, T]} |\epsilon_t^{(4)}(s)f| \leq \sup_{s \in [0, 1]} |\epsilon_{[t]}^{(4)}(s)f| + \sup_{s \in [0, T+1]} |\epsilon_{[t]}^{(4)}(s)f| + \sup_{s \in [0, 1]} |\epsilon_{[t+T]}^{(4)}(s)f| \leq \delta.$$

Since $\delta > 0$ is arbitrary, it follows that almost surely for every T

$$\lim_{t \rightarrow \infty} \sup_{s \in [0, T]} |\epsilon_t^{(4)}(s)f| = 0.$$

4.3.4. $\epsilon_t^{(5)}(s)f$

The final error term is a jump martingale term. We use the same approach as above for $\epsilon_t^{(4)}(s)f$. For this jump martingale the quadratic variation is

$$\sum_{0 < u \leq t} \left(Q_{\mu_u} f(X_u) \alpha_u - Q_{\mu_u} f(X_{u-}) \alpha_u \right)^2,$$

where the sum is over jump-points $u \in (0, t]$. The squared jump at time u is bounded by

$$16\|f\|_\infty^2 \kappa^{-2} \alpha_u^2,$$

making use of Lemma 4.3. The expected quadratic variation is then the expectation of the predictable variation, which is bounded by

$$\mathbb{E}\left[\int_0^t \kappa(X_u) \alpha_u^2 du\right] \leq \bar{\kappa} \int_0^t \alpha_u^2 du. \quad (40)$$

Thus the total quadratic variation is bounded, and as before we conclude that almost surely

$$\lim_{t \rightarrow \infty} \sup_{s \in [0, T]} |\epsilon_t^{(5)}(s)f| = 0.$$

4.3.5. Concluding the proof

We have shown that the five discrepancy terms all converge to 0 uniformly on compact sets as $t \rightarrow \infty$ almost surely. Thus condition (27) holds almost surely and the proof is complete. \square

4.4. Proof of Corollary 1.3

This follows from Theorem 1.1 and Remark 3.7, since limit sets of asymptotic pseudo-trajectories are *attractor free sets*, and so will be contained in the global attractor of Φ , which is $\{\pi\}$. These relationships are spelled out in [2, Section 5]. \square

5. Some comments on noncompact manifolds

In this work we have restricted our analysis to compact state spaces. Practically speaking, though, QSMC methods such as ScaLE [23] and ReScaLE as described here are applicable on noncompact state spaces such as \mathbb{R}^d .

The key difficulties in extending this present work to the setting of noncompact state spaces are:

1. establishing almost sure tightness of the occupation measures, and
2. arguing that $\mathbb{E}[\int_0^t \kappa(X_s) ds] = O(t)$ as $t \rightarrow \infty$.

We have also implicitly assumed there is a lower bound on the Ricci curvature, for the local bounds on the growth of the gradient of the Brownian transition kernel used in Section 4.3.3.

To establish tightness it might be helpful to consider the discrete skeleton at the regeneration times, $(\mu_{T_n})_{n \in \mathbb{N}}$, as a *measure-valued Pólya urn process*, as introduced in the recent work [22].

The second point – bounding $\mathbb{E}[\int_0^t \kappa(X_s) ds]$ – is necessary in controlling the variance of the jump martingale in (40).

[20] works in a noncompact setting and considers a related problem, namely studying self-interacting diffusions where the occupation measures (μ_t) come into play through the drift of the diffusion. However the techniques utilized – such as the construction of explicit Lyapunov functions – are not immediately applicable. Our choice of the metric (5) also needs to be modified, since – contrary to what is stated in [20, Section 2.1.1] – the space of bounded continuous functions is in general not separable on a noncompact state space.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Research of A.Q. Wang is supported by EPSRC OxWaSP CDT through grant EP/L016710/1. Research of G.O. Roberts is supported by EPSRC grants EP/K034154/1, EP/K014463/1, EP/D002060/1. We would particularly like to thank the anonymous referees whose comments have substantially improved the paper.

References

- [1] D. Aldous, B. Flannery, J.L. Palacios, Two applications of Urn processes: The fringe analysis of search trees and the simulation of quasi-stationary distributions of Markov chains, *Probab. Engrg. Inform. Sci.* 2 (03) (1988) 293–307, URL: <http://digitalassets.lib.berkeley.edu/sdtr/ucb/text/122.pdf>.
- [2] M. Benaïm, Dynamics of stochastic approximation algorithms, in: *Séminaire de Probabilités, XXXIII*, in: *Lecture Notes in Math.*, vol. 1709, Springer, Berlin, 1999, pp. 1–68, <http://dx.doi.org/10.1007/BFb0096509>.
- [3] M. Benaïm, B. Cloez, A stochastic approximation approach to quasi-stationary distributions on finite spaces, *Electron. Commun. Probab.* 20 (13) (2015) 1–14, <http://dx.doi.org/10.1214/ECP.v20-3956>.
- [4] M. Benaïm, B. Cloez, F. Panloup, Stochastic approximation of quasi-stationary distributions on compact spaces and applications, *Ann. Appl. Probab.* 28 (4) (2018) 2370–2416, <http://dx.doi.org/10.1214/17-AAP1360>.
- [5] M. Benaïm, M. Ledoux, O. Raimond, Self-interacting diffusions, *Probab. Theory Related Fields* 122 (1) (2002) 1–41, <http://dx.doi.org/10.1007/s004400100161>.
- [6] A. Beskos, O. Papaspiliopoulos, G.O. Roberts, Retrospective exact simulation of diffusion sample paths with applications, *Bernoulli* 12 (6) (2006) 1077–1098, <http://dx.doi.org/10.3150/bj/1165269151>.
- [7] J.-M. Bismut, Large Deviations and the Malliavin Calculus, in: *Progress in Mathematics*, vol. 45, Birkhäuser, 1984, p. 216.
- [8] J. Blanchet, P. Glynn, S. Zheng, Analysis of a stochastic approximation algorithm for computing quasi-stationary distributions, *Adv. Appl. Probab.* 48 (3) (2016) 792–811, <http://dx.doi.org/10.1017/apr.2016.28>.
- [9] N. Champagnat, D. Villemonais, Exponential convergence to quasi-stationary distribution and Q-process, *Probab. Theory Related Fields* 164 (1–2) (2016) 243–283, <http://dx.doi.org/10.1007/s00440-014-0611-7>, URL: <http://link.springer.com/10.1007/s00440-014-0611-7> arXiv:1506.02385.
- [10] S.N. Cohen, R.J. Elliott, *Stochastic calculus and applications*, second ed., Probability and its Applications, Springer, Cham, 2015, p. xxiii+666, <http://dx.doi.org/10.1007/978-1-4939-2867-5>.
- [11] P. Collet, S. Martínez, J. San Martín, Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems, Probability and its Applications (New York), Springer, Heidelberg, 2013, p. xvi+280, <http://dx.doi.org/10.1007/978-3-642-33131-2>.
- [12] M. Demuth, J.A. Casteren, *Stochastic Spectral Theory for Selfadjoint Feller Operators: a functional integration approach*, Birkhäuser Basel, 2000, p. 463.
- [13] L. Devroye, *Non-Uniform Random Variate Generation*, Springer-Verlag, 1986, p. 843.
- [14] D. Down, S.P. Meyn, R.L. Tweedie, Exponential and uniform ergodicity of Markov processes, *Ann. Probab.* 23 (4) (1995) 1671–1691, URL: [http://links.jstor.org/sici?sici=0091-1798\(199510\)23:4<1671:EAUEOM>2.0.CO;2-7&origin=MSN](http://links.jstor.org/sici?sici=0091-1798(199510)23:4<1671:EAUEOM>2.0.CO;2-7&origin=MSN).
- [15] R.M. Dudley, *Real Analysis and Probability*, in: *Cambridge Studies in Advanced Mathematics*, vol. 74, Cambridge University Press, Cambridge, 2002, p. x+555, <http://dx.doi.org/10.1017/CBO9780511755347>, Revised reprint of the 1989 original.
- [16] K.D. Elworthy, *Stochastic Differential Equations on Manifolds*, in: *London Mathematical Society Lecture Note Series*, vol. 70, Cambridge University Press, Cambridge, 1982, <http://dx.doi.org/10.1017/CBO9781107325609>.
- [17] S.N. Ethier, T.G. Kurtz, *Markov Processes: Characterization and Convergence*, in: *Wiley Series in Probability and Statistics*, Wiley, ISBN: 9780470316658, 1986, p. 534.
- [18] A. Guionnet, B. Zegarlinski, Lectures on logarithmic Sobolev inequalities, in: *Séminaire de Probabilités (Strasbourg)*, Vol. 36, 2002, pp. 1–134, <http://dx.doi.org/10.2307/2373688>.
- [19] M. Kolb, D. Steinsaltz, Quasilimiting behavior for one-dimensional diffusions with killing, *Ann. Probab.* 40 (1) (2012) 162–212, <http://dx.doi.org/10.1214/10-AOP623>.
- [20] A. Kurtzmann, The ODE method for some self-interacting diffusions on \mathbb{R}^d , *Ann. Inst. Henri Poincaré Probab. Stat.* 46 (3) (2010) 618–643, <http://dx.doi.org/10.1214/09-AIHP206>.

- [21] H.J. Kushner, G.G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, second ed., *Applications of Mathematics* (New York), vol. 35, Springer-Verlag, New York, 2003, p. xxii+474, *Stochastic Modelling and Applied Probability*.
- [22] C. Mailler, J.-F. Marckert, Measure-valued Pólya urn processes, *Electron. J. Probab.* 22 (26) (2017) 1–33, <http://dx.doi.org/10.1214/17-EJP47>.
- [23] M. Pollock, P. Fearnhead, A.M. Johansen, G.O. Roberts, The Scalable Langevin Exact Algorithm: Bayesian inference for big data, 2016. URL: <https://arxiv.org/abs/1609.03436>.
- [24] D. Revuz, M. Yor, *Continuous martingales and brownian motion*, third ed., *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, 293, Springer-Verlag, Berlin, 1999, p. xiv+602, <http://dx.doi.org/10.1007/978-3-662-06400-9>.
- [25] G.O. Roberts, J.S. Rosenthal, General state space Markov chains and MCMC algorithms, *Probab. Surv.* 1 (2004) 20–71, <http://dx.doi.org/10.1214/1549578041000000024>.
- [26] S. Schreiber, J.M. Levine, O. Godoy, N. Kraft, S. Hart, Does deterministic coexistence theory matter in a finite world? Insights from serpentine annual plants, *bioRxiv preprint*, 2018. <http://dx.doi.org/10.1101/290882>.
- [27] D.W. Stroock, *An Introduction to the Analysis of Paths on a Riemannian Manifold*, in: *Mathematical Surveys and Monographs*, vol. 74, American Mathematical Society, Providence, RI, 2000, p. xviii+269.
- [28] A.Q. Wang, M. Kolb, G.O. Roberts, D. Steinsaltz, Theoretical properties of quasi-stationary Monte Carlo methods, *Ann. Appl. Probab.* 29 (1) (2019) 434–457, <http://dx.doi.org/10.1214/18-AAP1422>.