

REGENERATION-ENRICHED MARKOV PROCESSES WITH APPLICATION TO MONTE CARLO

BY ANDI Q. WANG¹, MURRAY POLLOCK², GARETH O. ROBERTS³ AND DAVID STEINSALTZ⁴

¹*Department of Mathematics, University of Bristol, andi.wang@bristol.ac.uk*

²*School of Mathematics, Statistics and Physics, Newcastle University, Murray.Pollock@newcastle.ac.uk*

³*Department of Statistics, University of Warwick, gareth.o.roberts@warwick.ac.uk*

⁴*Department of Statistics, University of Oxford, steinsal@stats.ox.ac.uk*

We study a class of Markov processes that combine local dynamics, arising from a fixed Markov process, with regenerations arising at a state-dependent rate. We give conditions under which such processes possess a given target distribution as their invariant measures, thus making them amenable for use within Monte Carlo methodologies. Since the regeneration mechanism can compensate the choice of local dynamics, while retaining the same invariant distribution, great flexibility can be achieved in selecting local dynamics, and the mathematical analysis is simplified. We give straightforward conditions for the process to possess a central limit theorem, and additional conditions for uniform ergodicity and for a coupling from the past construction to hold, enabling exact sampling from the invariant distribution. We further consider and analyse a natural approximation of the process which may arise in the practical simulation of some classes of continuous-time dynamics.

1. Introduction. In this work we study a broad class of continuous-time Markov processes X which are defined by superimposing regenerative dynamics onto an existing continuous-time Markov process Y on state space E . The precise definition of the process is given in Section 2, but X can be seen informally as a Markov process with infinitesimal generator L^μ given by

$$(1) \quad L^\mu f(x) = Qf(x) + \kappa(x) \int (f(y) - f(x))\mu(y)m(dy),$$

where Q denotes the infinitesimal generator of the process Y .

We will refer to the function κ as the *regeneration rate* and to μ as the *regeneration density*. Collectively, κ and μ constitute the *global dynamics*. We will refer to the dynamics defined by the process Y as the *local dynamics*. For example, we may choose the local dynamics to be a Brownian motion on \mathbb{R}^d or a continuous-time jump process, such as a suitably-defined Metropolis–Hastings chain embedded in continuous time.

Fundamental to the introduction of this class of Markov processes—which we term *Restore* processes as they are *Randomly Exploring and STOchastically REgenerating*—is that the global regenerative dynamics we introduce can enrich the existing local dynamics of Y in a compensatory manner, such that the invariant distribution of X is known. This is of particular application within Monte Carlo methodology as the dynamics can often be chosen, and the process straightforwardly simulated, such that its invariant distribution coincides with a prescribed target density of interest, π .

Received October 2019; revised June 2020.

MSC2020 subject classifications. Primary 60J40, 60J22; secondary 60J25, 65C05.

Key words and phrases. Right process, inhomogeneous Poisson process, regenerative Markov process, Markov chain Monte Carlo, coupling from the past.

This incorporation of global regenerative dynamics to enrich an existing continuous-time Markov process Y introduces a number of directions for both theoretical and methodological innovation, which we explore in this paper. Since the process naturally exhibits regenerations, mathematical analysis of the Restore process is simplified considerably. Indeed, the traditional approach of analysing Markov chains, in the sense of [Meyn and Tweedie \(1993\)](#), crucially relies on the identification of regeneration times.

Thus the Restore process forms the basis of a new approach to Monte Carlo sampling, which we term the *Restore sampler*. The Restore sampler also provides a simple recipe for introducing (nonreversible) rejection-free moves to existing samplers. This can be done in cases where standard Markov chain Monte Carlo (MCMC) algorithms may exhibit poor mixing. This is discussed in [Sections 3.2](#).

The Restore process is an instance of a “resurrected” or “returned” process, which instantaneously returns to the state space after being killed. Such processes have been utilized extensively within probability literature. Their use goes back to the very foundations of Markov chain theory, [Doob \(1945\)](#), but such processes have been harnessed particularly effectively in the study of *quasi-stationarity*. See, for instance, [Bartlett \(\(1960\), Section 3.4\)](#), [Darroch and Seneta \(1965\)](#), [Collet, Martínez and San Martín \(\(2013\), Chapter 4.4\)](#), [Barbour and Pollett \(2010, 2012\)](#), [Benaim, Cloez and Panloup \(2018\)](#), [Wang, Roberts and Steinsaltz \(2020\)](#), [Wang and Steinsaltz \(2019\)](#). For example, such processes have been used to approximate quasi-stationary distributions, and in this context, [Darroch and Seneta \(1965\)](#) noted that for discrete-time, finite state space resurrected processes, the invariant distribution could be “made into almost any distribution”. The work of this paper demonstrates that for continuous time and general state spaces, this is also the case.

The idea of identifying regeneration times within a given MCMC sampler goes back to [Mykland, Tierney and Yu \(1995\)](#), using the very elegant splitting technique of [Nummelin \(1978\)](#). The area has continued to develop actively, as seen, for instance, in the contributions of [Brockwell and Kadane \(2005\)](#), [Gilks, Roberts and Sahu \(1998\)](#), [Hobert et al. \(2002\)](#), [Lee, Doucet and Latuszyński \(2014\)](#), [Minh, Minh and Nguyen \(2012\)](#). The idea of hybridising separate dynamics has also had a long history, see, for instance, [Murdoch \(2000\)](#), [Murdoch and Green \(1998\)](#), [Tierney \(1996\)](#), although these typically involve combining separate MCMC chains which are already themselves π -invariant. The Restore process offers practitioners considerable scope to design highly optimised sampling algorithms due to the flexibility of being able to “hybridise” dynamics which are separately not π -invariant.

Unlike traditional MCMC methods, the Restore sampler is a fundamentally continuous-time sampler, as the inhomogeneous Poisson clock dictating the regeneration events is crucial for aligning the local and global dynamics. In a similar vein, the class of *piecewise-deterministic Markov processes* (PDMPs, [Davis \(1984\)](#)), and *quasi-stationary Monte Carlo methods* (QSMC) also make use of an inhomogeneous Poisson process to drive the process towards the target distribution π ; see [Vanetti et al. \(2017\)](#), [Wang et al. \(2019\)](#). Notable examples of such methods include the Bouncy Particle Sampler, [Bouchard-Côté, Vollmer and Doucet \(2018\)](#), the Zig-zag Sampler, [Bierkens, Fearnhead and Roberts \(2019\)](#), ScaLE, [Pollock et al. \(2020\)](#) and ReScaLE, [Kumar \(2019\)](#).

Sampling algorithms which rely upon continuous-time dynamics often require some form of approximation for their practical implementation, and the resulting approximate process can exhibit algorithmic instability, or possess an approximate invariant distribution which is intractable. However, there is considerable scope and promise to understand the effect of such approximations with the Restore process, due to the global regenerative dynamics with which it is constructed and the ease with which it can be mathematically analysed (for instance, in the sense of [Asmussen and Glynn \(2007\)](#)). In [Section 5.2](#), we consider one natural approximation to the Restore process in which the regeneration rate is truncated.

1.1. *Summary of results.* We begin in Section 2 by formally introducing the Restore process on an abstract state space E , and in Section 3 we will establish the following.

CONTRIBUTION 1 (π -invariance: Theorems 16, 22). Assume that we are given a positive target density π on E , a regeneration density μ with on E , and an interarrival process Y with infinitesimal generator Q with adjoint Q^* . We assume that we have chosen a constant $C > 0$ such that

$$(2) \quad Q^* \pi + C \mu \geq 0.$$

Under a range of settings and regularity conditions, to be detailed in Section 3, the resulting Restore process with interarrival dynamics Y , regeneration rate κ and regeneration density μ has invariant density π .

We consider the following two indicative settings: symmetric diffusion processes and continuous-time jump processes.

In Section 4, we study limiting properties of the Restore process and will present the following results.

CONTRIBUTION 2 (Central limit theorem: Theorem 24). Writing (T_n) for the regeneration times, then for appropriate functions f , where σ_f^2 is the asymptotic variance defined in (18), then under appropriate regularity conditions the following holds for the Restore process:

$$\sqrt{n} \left(\frac{\int_0^{T_n} f(X_s) ds}{T_n} - \pi[f] \right) \xrightarrow{d} N(0, \sigma_f^2).$$

Under additional assumptions, we will derive uniform ergodicity and a *coupling from the past* (CFTP) construction (following [Propp and Wilson \(1996\)](#)), which is particularly useful in the context of Monte Carlo simulation, since it allows us to obtain an exact draw from the target π .

CONTRIBUTION 3 (Uniform ergodicity, CFTP: Proposition 26, Theorem 27). Assume that the regeneration rate κ is uniformly bounded away from 0 and basic regularity conditions hold. Then the Restore process is uniformly ergodic. Furthermore, there is a straightforward coupling from the past construction.

Indeed, in Theorem 28, we show that the classical rejection sampler is a special case of this coupling from the past construction.

In Section 5, we discuss some practical considerations related to the Restore sampler, and in particular we present a result concerning the error incurred when running one natural approximation of the Restore process.

CONTRIBUTION 4 (Truncated rate: Theorem 30, Proposition 32). When the interarrival process is a diffusion and κ is bounded away from 0, consider running the Restore process with a *truncated* version of the regeneration rate κ_M :

$$\kappa_M := \kappa \wedge M.$$

Writing π_M for the invariant distribution of the resulting approximate process, we provide a bound on the error $\|\pi_M - \pi\|_1$ in total variation and show it vanishes to 0 as the truncation level $M \rightarrow \infty$.

Some simple examples highlighting various aspects of the Restore sampler are given in Section 6. To conclude, in Section 7 we discuss the limitations of our approach, and possible future directions. Some technical proofs are omitted from the body of the text for readability, but can found in the [Appendix](#).

2. The Restore process. First we formalize the informal definition of the Restore process given in the [Introduction](#). We define the process in a general, abstract framework.

Let (E, m) be a measure space, where E is a Radon topological space with its Borel σ -algebra \mathcal{E} and m is a σ -finite Radon measure on \mathcal{E} , for example, \mathbb{R}^d equipped with Lebesgue measure. We assume that we are given a *right process* $Y = (\Omega', \mathcal{F}', \mathcal{F}'_t, Y_t, \mathbb{P}_x^0)$ evolving on E . Right processes are an abstract class of right-continuous strong Markov processes. We do not repeat their precise definition, which is highly technical, here; the interested reader is referred to [Sharpe \(\(1988\), Chapter 20\)](#), instead we give a list of examples in the following lemma.

LEMMA 1. *The following processes are examples of right processes: deterministic right-continuous flows, Feller processes, Markov jump processes.*

PROOF. See [Sharpe \(\(1988\), Exercises 8.8, 9.27 and 14.18\)](#). \square

REMARK 2. Recall that a Feller process is a Markov process on a locally compact, Hausdorff, second countable space E , whose semigroup (P_t) is strongly continuous on $C_0(E)$, the set of continuous functions vanishing at infinity. Examples of Feller processes include Lévy processes, [Sharpe \(\(1988\), page 50\)](#), and diffusions such as the ones studied in [Demuth and van Casteren \(\(2000\), Chapter 1\)](#).

For a general initial distribution ν we write $\mathbb{P}_\nu^0 = \int_E \nu(dx) \mathbb{P}_x^0$. Let $\kappa : E \rightarrow \mathbb{R}^+ = [0, \infty)$ be a locally bounded measurable function, the *regeneration rate*. Define the *lifetime* τ_∂ as

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

where $\xi \sim \text{Exp}(1)$, independent of Y . Set $\inf \emptyset = \infty$.

Fix a probability measure ν_μ on (E, m) , the *regeneration distribution*. We define the Restore process $X = (X_t)_{t \geq 0}$ to be the process given by

$$(4) \quad X_t = \sum_{i=0}^{\infty} 1_{[T_i, T_{i+1})}(t) Y_{t-T_i}^{(i)},$$

where $(Y^{(0)}, \tau^{(0)})$ is a realisation of (Y, τ_∂) with $Y_0 = x$, and $(Y^{(i)}, \tau^{(i)})_{i=1}^\infty$ are i.i.d. realisations of (Y, τ_∂) under $\mathbb{P}_{\nu_\mu}^0$, namely with $Y_0 \sim \nu_\mu$. The $(T_i)_{i=0}^\infty$ are given by $T_0 = 0$, and $T_n = \sum_{i=0}^{n-1} \tau^{(i)}$, for each $n = 1, 2, \dots$.

This defines a Markov process $X = (\Omega, \mathcal{F}, \mathcal{F}_t, X_t, \mathbb{P}_x)$ with state space (E, m) . For an arbitrary initial distribution ν , as usual we set $\mathbb{P}_\nu = \int \nu(dx) \mathbb{P}_x$. In future, the regeneration measure ν_μ will be given by a density function μ with respect to the reference measure m , and hence for its semigroup we will write $\{P_t^\mu : t \geq 0\}$. We will then refer to this process as the *Restore process with interarrival dynamics Y , regeneration rate κ , and regeneration density μ* .

LEMMA 3. *Let Y be a right process on the Radon space (E, m) with Radon measure m , $\kappa : E \rightarrow \mathbb{R}^+$ a locally bounded measurable function, and μ a probability measure on E . Then the resulting Restore process $X = (\Omega, \mathcal{F}, \mathcal{F}_t, X_t, \mathbb{P}_x)$ with interarrival dynamics Y , locally bounded nonnegative regeneration rate κ and regeneration density μ defines a right process with state space (E, m) . In particular, X is right-continuous and strong Markov. Moreover, $T_n \rightarrow \infty$ almost surely.*

PROOF. See [Appendix A.1](#). \square

3. Invariance. Suppose we are given a probability measure ν_π on E , our target measure of interest. We will assume throughout that the target measure ν_π and regeneration measure ν_μ are given by density functions π, μ respectively with respect to the reference measure m , namely

$$\nu_\pi(dx) = \pi(x)m(dx), \quad \nu_\mu(dx) = \mu(x)m(dx).$$

We would like to construct a Restore process X whose invariant distribution coincides with π . In this section we formulate conditions in several settings under which this is possible.

We consider the following settings: when the interarrival process Y is a symmetric diffusion, and when the interarrival process is a jump process. This latter situation includes, for example, the case when the state space E is countable.

Writing Q for the generator of the process Y , define the regeneration rate $\kappa : E \rightarrow \mathbb{R}$ by

$$(5) \quad \kappa(x) := \frac{Q^*\pi(x)}{\pi(x)} + C \frac{\mu(x)}{\pi(x)}, \quad x \in E.$$

We will make rigorous sense of this expression in the subsequent sections.

REMARK 4. Because of the flexibility provided by the constant C in (5), in practice we do not require μ or π to be normalized in order to compute κ .

Given the formal generator (1), we can make intuitive sense of the expression (5) from the following formal manipulations:

$$\nu_\pi Qf = \int \pi(x)(Qf)(x)m(dx) = \int (Q^*\pi)(x)f(x)m(dx).$$

Taking $f \equiv 1$ the constant function, we see that $\int Q^*\pi(x)m(dx) = 0$, since $Q1 \equiv 0$. Then,

$$\begin{aligned} \nu_\pi[\kappa(\nu_\mu[f] - f)] &= \nu_\pi[\pi^{-1}(Q^*\pi + C\mu)]\nu_\mu[f] - \nu_\pi[\pi^{-1}fQ^*\pi] - C\nu_\pi[\pi^{-1}\mu f] \\ &= 0 - m[fQ^*\pi] + C(m[\mu f]m[\mu] - m[\mu f]). \end{aligned}$$

This final bracket is 0 since $m[\mu] = 1$, as ν_μ is a probability measure. This allows us to conclude that

$$\nu_\pi[L^\mu f] = \nu_\pi[Qf] - \nu_\pi[\kappa(\nu_\mu[f] - f)] = m[\pi Qf] - m[fQ^*\pi] = 0.$$

This calculation shows that our κ is indeed of the right form to ensure invariance of ν_π .

We emphasize again that the preceding calculations are formal and do not constitute a rigorous proof. In order to turn this into a full proof, one must first show that the operator L^μ given in (1) is indeed the generator of the Restore process (as constructed in Section 2), carefully noting the domain $\mathcal{D}(L^\mu)$. We must then establish that the above calculations hold for a collection of functions $f \in D$, and prove that D constitutes a core of the generator.

REMARK 5. Turning these calculations into a proof in a general setting is difficult for several reasons. First, establishing that L^μ is the generator of the Restore process is complicated since κ is not necessarily bounded, thus the Restore process is not necessarily Feller in the sense of Remark 2. This prevents us from straightforwardly establishing dissipativity, via the positive maximum principle, which would enable the application of general reformulations of the Hille–Yosida theorem such as Theorem 7.1 of Ethier and Kurtz (1986). Second, proving that a collection of functions D constitute a core for the generator is generally challenging. For recent advances on this topic for PDMPs, see the work of Durmus, Guillin and Monmarché (2018).

These difficulties associated with working in a general operator-theoretic setting are our motivation for considering our two specific settings separately; the diffusion setting, and the jump process setting. Indeed, one of the key contributions of this work is that in each setting we will give a proof of invariance which avoids using the full generator approach and the highly technical difficulties outlined above.

Once invariance is established, in order to approximate integrals we can make use of the following result. Recall that a nonnegative random variable is *nonlattice* if it is not concentrated on a set of the form $\{\delta, 2\delta, \dots\}$ for any $\delta > 0$.

THEOREM 6. *Suppose that the Restore process X , as in the conclusion of Lemma 3, is defined on a metric space E , its semigroup P_t^μ maps continuous functions to continuous functions for each $t \geq 0$, has a unique stationary distribution π , that $\mathbb{E}_\mu[\tau_\delta] < \infty$, and that the lifetimes are nonlattice. Then for any bounded measurable function $f : E \rightarrow \mathbb{R}$, we have that*

$$(6) \quad \nu_\pi[f] = \frac{\mathbb{E}_\mu[\int_0^{\tau^{(0)}} f(X_s) ds]}{\mathbb{E}_\mu[\tau^{(0)}]},$$

and furthermore we have almost sure convergence of the ergodic averages: as $t \rightarrow \infty$,

$$\frac{1}{t} \int_0^t f(X_s) ds \rightarrow \nu_\pi[f].$$

PROOF. By Theorem 1.2 of [Asmussen \(\(2003\), Chapter 6\)](#), and uniqueness of the stationary distribution, it follows that (6) holds. Convergence of the ergodic averages then follows from the following arguments from renewal theory: First split f into positive and negative parts, so we may assume that f is nonnegative. Writing $(N(t))_{t \geq 0}$ for the renewal process of complete lifetimes before time t , we may thus bound

$$\int_0^{T_{N(t)}} f(X_s) ds \leq \int_0^t f(X_s) ds \leq \int_0^{T_{N(t)+1}} f(X_s) ds.$$

By the strong law of large numbers for renewal processes, Theorem 1 of [Grimmett and Stirzaker \(\(2001\), 10.2\)](#), we know that $N(t)/t \rightarrow 1/\mathbb{E}_\mu[\tau^{(0)}]$ almost surely. We can conclude the argument by then applying the strong law of large numbers to $\int_0^{T_{N(t)}} f(X_s) ds/N(t)$ and similarly for the upper bound. \square

3.1. Symmetric diffusions. We first consider Restore when the underlying process is a symmetric diffusion on $E = \mathbb{R}^d$. For a smooth C^∞ function $A : \mathbb{R}^d \rightarrow \mathbb{R}$ consider the stochastic differential equation (SDE)

$$(7) \quad dY_t = \nabla A(Y_t) dt + dB_t, \quad Y_0 = x,$$

on \mathbb{R}^d where B is a standard Brownian motion on \mathbb{R}^d . Define the smooth function $\gamma : \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$\gamma(y) = \exp(2A(y)), \quad y \in \mathbb{R}^d,$$

and define a measure Γ on \mathbb{R}^d by

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

where dy denotes Lebesgue measure on \mathbb{R}^d .

We are thus working on $(E, m) = (\mathbb{R}^d, \Gamma)$. This is an example of a Radon space with a Radon measure.

ASSUMPTION 7 (Underlying process). $A : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth C^∞ function, and the SDE (7) has a unique weak solution. The process Y has a continuous symmetric transition density $p^0(t, x, y)$ on $(0, \infty) \times \mathbb{R}^d \times \mathbb{R}^d$ with respect to Γ , which satisfies the BASSA conditions of Demuth and van Casteren ((2000), Chapter 1.B). In particular, the diffusion is Feller, hence a right process.

The BASSA conditions of Demuth and van Casteren ((2000), Chapter 1.B) are technical, and in Section 3.1.1 we will give examples of diffusions satisfying them.

The semigroup of the diffusion Y is given for each $t \geq 0$ by

$$(8) \quad \mathbb{E}_x^0[f(Y_t)] = \int p^0(t, x, y) f(y) d\Gamma(y),$$

for functions f where this integral makes sense. Under Assumption 7, the semigroup (8) maps $C_0(\mathbb{R}^d)$ —continuous functions vanishing at ∞ —into $C_0(\mathbb{R}^d)$ and is strongly continuous on $C_0(\mathbb{R}^d)$ with generator $-L^0$. Hence we can also write the semigroup as

$$\mathbb{E}_x^0[f(Y_t)] = [\exp(-tL^0)f](x).$$

The action of the generator on smooth compactly supported f is given by

$$-L^0 f = \frac{1}{2} \Delta f + \nabla A \cdot \nabla f.$$

Note that we are writing L^0 for *minus* the generator, as is done in Demuth and van Casteren (2000).

Under Assumption 7, the semigroup is also strongly continuous on

$$\mathcal{L}^p(\Gamma) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ measurable, } \int_{\mathbb{R}^d} |f(x)|^p d\Gamma(x) < \infty \right\},$$

for each $1 \leq p < \infty$. When we want to emphasize the underlying function space we may write $-L_p^0$ for the corresponding generators on $\mathcal{L}^p(\Gamma)$ and $\mathcal{D}(L_p^0) \subset \mathcal{L}^p(\Gamma)$ for their respective dense domains.

We now assume that the target distribution and regeneration distributions are defined by density functions with respect to Γ denoted $\pi, \mu \in \mathcal{L}^1(\Gamma)$ respectively.

ASSUMPTION 8 (Densities). The target density $\pi \in \mathcal{L}^1(\Gamma)$, is positive on \mathbb{R}^d and is twice continuously differentiable with $\int \pi d\Gamma = 1$. The regeneration density μ is in $\mathcal{L}^1(\Gamma)$ and is nonnegative, with $\int \mu d\Gamma = 1$. Furthermore, π and μ are square-integrable—that is, in $\mathcal{L}^2(\Gamma)$ —and π is in the domain $\mathcal{D}(L_2^0)$.

REMARK 9. Let us emphasize that we are writing π and μ for densities with respect to the measure Γ , which may not necessarily be Lebesgue measure. Later on we will write $\bar{\pi} := \pi \gamma$ for the density with respect to Lebesgue measure.

For our proofs we take π, μ to be normalized, but as noted previously this condition is not required in practice, because of the constant C which appears in the regeneration rate.

Because L^0 is a self-adjoint operator on $\mathcal{L}^2(\Gamma)$, a sufficient condition for $\pi \in \mathcal{D}(L_2^0)$ is that $L^0 \pi \in \mathcal{L}^2(\Gamma)$. This is a well-known result; for a derivation, see, for example, Wang ((2020), Section 3.3.3), where a preliminary version of this work can also be found.

We can now define the regeneration rate κ , under Assumption 8. First, define the *partial regeneration rate* $\tilde{\kappa}$, via

$$\tilde{\kappa}(x) := \frac{1}{\pi(x)} \left(\frac{1}{2} \Delta \pi(x) + \nabla A \cdot \nabla \pi(x) \right), \quad x \in \mathbb{R}^d.$$

We define the actual regeneration rate κ as follows. Set for a given constant $C > 0$,

$$(9) \quad \kappa(x) := \tilde{\kappa}(x) + C \frac{\mu(x)}{\pi(x)}, \quad x \in \mathbb{R}^d.$$

REMARK 10. Similar to Wang et al. (2019), writing $U := -\log \pi$, an equivalent expression for $\tilde{\kappa}$ is

$$(10) \quad \tilde{\kappa}(x) = \frac{1}{2}(-\Delta U(x) + |\nabla U(x)|^2) - \nabla A \cdot \nabla U(x).$$

ASSUMPTION 11 (Regeneration rate). The function κ is continuous, and C is chosen such that $\kappa \geq 0$.

Under Assumptions 7, 8, 11, the process Y killed at rate κ , that is, with lifetime given by (3), can be analysed using Theorem 2.5 of Demuth and van Casteren (2000).

PROPOSITION 12. Under Assumptions 7, 8, 11, the process Y killed at rate κ , that is, with lifetime given by (3), defines a strongly continuous sub-Markovian semigroup $\{\exp(-tL^\kappa) : t \geq 0\}$ on $C_0(\mathbb{R}^d)$ with symmetric, continuous kernel $p^\kappa(t, x, y)$. The corresponding generator $-L^\kappa = -L^0 - \kappa$, extends $-L^0 - \kappa$. In addition, it has Feynman–Kac representation,

$$\begin{aligned} [\exp(-tL^\kappa)f](x) &= \int p^\kappa(t, x, y) f(y) \, dm(y) \\ &= \mathbb{E}_x \left[\exp\left(-\int_0^t \kappa(Y_s) \, ds\right) f(Y_t) \right]. \end{aligned}$$

Furthermore, the semigroup is strongly continuous on $\mathcal{L}^p(\Gamma)$ for any $1 \leq p < \infty$. In particular, on $\mathcal{L}^2(\Gamma)$, it is self-adjoint and possesses a self-adjoint generator.

PROOF. See Appendix A.2. \square

As before, when we want to make explicit which $\mathcal{L}^p(\Gamma)$ space we are using, for $1 \leq p < \infty$, we will write $-L_p^\kappa$ for the generator of the strongly continuous semigroup on $\mathcal{L}^p(\Gamma)$, with corresponding domain $\mathcal{D}(L_p^\kappa) \subset \mathcal{L}^p(\Gamma)$. The domain of the generator may be defined as the image of the semigroup acting on $\mathcal{L}^p(\Gamma)$.

REMARK 13. It follows from Assumption 8 that $\pi \in \mathcal{D}(L_2^\kappa)$, since both π and μ are in $\mathcal{L}^2(\Gamma)$, and formally $L^\kappa \pi = C\mu$.

We have one final technical assumption.

ASSUMPTION 14 (Technical conditions on π, μ). We have that

$$(11) \quad \pi \in \mathcal{D}(L_1^\kappa), \quad L_1^\kappa \pi = C\mu.$$

Furthermore, μ is such that

$$(12) \quad \int d\Gamma(x) \mu(x) \mathbb{E}_x^0 \left[\sup_{t \in [0,1]} |\kappa(Y_t) e^{-\int_0^t \kappa(Y_s) \, ds}| \right] < \infty.$$

The condition (11) is fairly abstract, and so might be difficult to verify in a particular case, or in a general class of processes that one may want to consider. Lemma 15 gives a sufficient condition which we will make use of. Set

$$\bar{\pi} := \pi \gamma.$$

We write $W^{2,1}(\mathbb{R}^d)$ for the Sobolev space of measurable functions on \mathbb{R}^d whose first and second derivatives are integrable with respect to Lebesgue measure on \mathbb{R}^d .

LEMMA 15. *Assume that Assumptions 7, 8, 11 hold. Suppose that the drift is at most linear in the tails: we can bound $|\nabla A(x)| \leq K|x|$, for some $K > 0$, for all x outside of some compact set. Suppose π is smooth, and that $\bar{\pi} \in W^{2,1}(\mathbb{R}^d)$. In addition, we require that*

$$\int_{\mathbb{R}^d} |\nabla A(x) \cdot \nabla \bar{\pi}(x)| dx < \infty, \quad \int_{\mathbb{R}^d} |\Delta A(x) \bar{\pi}(x)| dx < \infty.$$

Then (11) holds.

PROOF. See Appendix A.3. \square

Alternatively, (11) will automatically hold whenever Γ is a finite measure. This is the case whenever the underlying diffusion Y is positive recurrent, say a stable Ornstein–Uhlenbeck process. Then under $\pi \in \mathcal{L}^2(\Gamma)$ and $\mu \in \mathcal{L}^2(\Gamma)$, $\pi \in \mathcal{D}(L_1^\kappa)$ with $L_1^\kappa \pi = C\mu$, since in that case \mathcal{L}^2 convergence implies \mathcal{L}^1 convergence.

The condition (12) is needed so that we can differentiate under the integral. A necessary condition for (12) to hold is that $\int d\Gamma(x)\mu(x)\kappa(x) < \infty$, so in particular μ cannot have tails which are too heavy relative to π . From a computational point of view, this is reasonable since otherwise the regeneration mechanism would be highly inefficient; the Restore process would tend to regenerate very rapidly. Of course, a sufficient condition for (12) is that

$$\int d\Gamma(x)\mu(x) \mathbb{E}_x^0 \left[\sup_{t \in [0,1]} \kappa(Y_t) \right] < \infty.$$

THEOREM 16. *Under Assumptions 7, 8, 11, 14, the Restore process X with interarrival dynamics Y , regeneration rate κ and regeneration density μ has invariant distribution π .*

PROOF. See Appendix A.4. \square

3.1.1. *Examples.* We now give some examples of diffusions which satisfy the assumptions of Theorem 16.

Sufficient conditions ensuring BASSA are given in Example 2 of Demuth and van Casteren ((2000), Chapter 1.C). In our present setting when we consider diffusions defined by (7), these conditions can be written as

$$(13) \quad \exp(A(x)) \geq c^{-1} \exp(-c|x|^2) \quad \forall x \in \mathbb{R}^d,$$

$$(14) \quad c^{-1} \leq \exp(A(x) - A(y)) \leq c \quad \forall x, y \in \mathbb{R}^d : |x - y| \leq c^{-1}(1 + |x|)^{-c},$$

for some $c > 0$.

Let $|\cdot|$ denote the ℓ_2 norm on \mathbb{R}^d .

PROPOSITION 17. *The SDE (7) with $A = \alpha|x|^2$ for any $\alpha \in \mathbb{R}$ satisfies BASSA.*

REMARK 18. In this case $\nabla A(x) = 2\alpha x$ is linear. $\alpha < 0$ corresponds to a (stable) Ornstein–Uhlenbeck process, $\alpha = 0$ is a Brownian motion and $\alpha > 0$ is an *unstable* Ornstein–Uhlenbeck process which drifts into the tails.

PROOF. (13) clearly holds in this setting. The second condition (14) can be seen from the reverse triangle inequality:

$$\begin{aligned} ||x|^2 - |y|^2| &= (|x| + |y|)|x| - |y|| \leq (|x| + |y|)|x - y| \\ &\leq \frac{|x| + |y|}{c(1 + |x|)^c} \leq \frac{2|x|}{c(1 + |x|)^c} + \frac{1}{c^2(1 + |x|)^{2c}}. \end{aligned}$$

This is uniformly bounded over $x \in \mathbb{R}^d$ for $c > 1$. \square

3.2. *Jump processes.* The Restore process is inherently a continuous-time process, and so the underlying process Y must be a continuous-time object. Suppose, however, we are given a discrete-time Markov transition kernel P on (E, m) , with action on measurable functions $f : E \rightarrow \mathbb{R}$ and measures ν on E given by

$$(15) \quad \begin{aligned} Pf(x) &= \int f(y)p(x, y) dm(y), \quad x \in E, \\ \nu P(dy) &= \int \nu(dx) p(x, y) dm(y), \end{aligned}$$

for some integral kernel $p(x, y)$ on $E \times E$, whenever these integrals make sense. Since we have an integral kernel $p(x, y)$, we will also think of νP as a measurable function given by

$$\nu P(y) := \int \nu(dx) p(x, y)$$

for a measure ν on E , provided this makes sense.

It is straightforward to embed P into continuous time, by specifying a measurable function $\lambda : E \rightarrow \mathbb{R}^+$, the *holding rates*. We take the jump chain to be defined by the discrete-time Markov kernel P , and just take the holding times to be independent $\text{Exp}(\lambda(x))$ times, when currently at state x .

Such a process will be a *continuous-time jump process* on E , meaning it has right-continuous, piecewise-constant sample paths. Provided they are nonexplosive, such processes are determined by the transition kernel of the jump chain and the holding rates. See, for instance, [Ethier and Kurtz \(\(1986\), Chapter 4.2\)](#).

Suppose π, μ are two densities on E with respect to m , the target density and regeneration density respectively, where we assume π is positive. Suppose we are given a transition kernel P on E and holding rates $\lambda : E \rightarrow \mathbb{R}^+$. We now construct the Restore process. Given a constant C , define the regeneration rate κ to be

$$(16) \quad \kappa(x) := \frac{\int \pi(y)\lambda(y)p(y, x)m(dy) - \lambda(x)\pi(x)}{\pi(x)} + C\frac{\mu(x)}{\pi(x)}, \quad x \in E.$$

ASSUMPTION 19 (Jump process Restore). P is a transition kernel with a density as in (15), $\lambda : E \rightarrow \mathbb{R}^+$ is measurable, strictly positive. π is a positive probability density with respect to m , μ is a probability density with respect to m , and $\int \lambda(x)\pi(x)m(dx) < \infty$. The constant C is such that $\kappa \geq 0$ on E . κ is locally bounded, and we have that $\int (\lambda(x) + \kappa(x))^2 \pi(x)m(dx) < \infty$.

Note that such jump processes are right processes (Exercise 14.18 of [Sharpe \(1988\)](#)).

REMARK 20. From (16), we see that a sufficient condition for $\kappa(x) \geq 0$ is that

$$C\mu(x) \geq \lambda(x)\pi(x), \quad x \in E.$$

Alternatively, if the underlying process is already π -invariant, so $\pi Q^0 \equiv 0$, then the first term in (16) is identically zero and any $C > 0$ and μ may be chosen.

REMARK 21. When the state space E is countable, the Markov process Y is necessarily a jump process, with law defined by a transition rate matrix Q . Probability distributions on E are given by (possibly infinite) row vectors. In this case, the regeneration rate (16) can be written simply as

$$\kappa(x) := \frac{(\pi Q)(x)}{\pi(x)} + C \frac{\mu(x)}{\pi(x)}, \quad x \in E,$$

where πQ is a row vector obtained by straightforward matrix multiplication.

This construction can also be extended to kernels P which do not possess a density as in (15). For example, the classical Metropolis–Hastings kernel is of the form

$$P(x, dy) = \alpha(x, y)q(x, y) dm(y) + (1 - j(x))\delta_x(dy),$$

where $0 \leq \alpha(x, y) \leq 1$ are the acceptance probabilities, $q(x, y)$ is a transition density (so $\int q(x, y) dm(y) = 1$ for each $x \in E$), and

$$j(x) := \int \alpha(x, y)q(x, y) dm(y)$$

are the jump probabilities. Because of the presence of the delta mass $\delta_x(dy)$, such kernels cannot possess straightforward densities. However in continuous-time, these rejected moves associated with the delta mass are not visible, and so we can modify the regeneration rate as follows: we replace the term $(\pi\lambda)P(x)$ in (16) by

$$\int dm(y) \pi(y)\lambda(y)\alpha(y, x)q(y, x) + \lambda(x)(1 - j(x))\pi(x).$$

Returning to the construction of the Restore process, we will take the interarrival dynamics to be given by the jump process defined by P and λ , the regeneration rate to be κ and the regeneration density μ . The resulting Restore process X is another continuous-time jump process, and so we describe its jump chain and holding rates. This will provide a method to simulate the process.

At $x \in E$, the transition kernel $P^\mu(x, dy)$ of the jump chain is given by

$$P^\mu(x, dy) = \frac{\lambda(x)}{\lambda(x) + \kappa(x)} P(x, dy) + \frac{\kappa(x)}{\lambda(x) + \kappa(x)} \mu(y) dm(y).$$

The overall holding rates in continuous time are given by

$$\bar{\lambda}(x) = \lambda(x) + \kappa(x), \quad x \in E,$$

that is, at $x \in E$, by the Markov property, the time until the next jump is an $\text{Exp}(\bar{\lambda}(x))$ time.

THEOREM 22. Assume that Assumption 19 holds, and that the interarrival dynamics defined by P and λ are nonexplosive. Then the resulting Restore process X is a nonexplosive jump process with invariant distribution π .

PROOF. Nonexplosivity follows from Lemma 3, and the fact that the interarrival process is assumed nonexplosive.

Let us write $\{Q_t^\mu : t \geq 0\}$ for the continuous-time semigroup for the Restore process X . Our goal is to show that $\pi Q_t^\mu f = \pi[f]$ for any continuous bounded function $f : E \rightarrow \mathbb{R}$, for each $t \geq 0$. To do this we compute the time derivative of the mapping $t \mapsto \pi Q_t^\mu f$, and show that it is 0. By time-homogeneity and the semigroup property, it is sufficient to compute this derivative at $t = 0$. This was the approach similarly used to prove π -invariance of

Algorithm 1 Jump process Restore Sampler

```

1: initialize:  $X_0 = x_0, t_0 = 0, i = 0$ 
2: while  $t_i < T$  do
3:    $i \leftarrow i + 1$ 
4:   simulate  $\tau_{i-1}^{(1)} \sim \text{Exp}(\lambda(X_{i-1})), \tau_{i-1}^{(2)} \sim \text{Exp}(\kappa(X_{i-1}))$ 
5:    $\tau_{i-1} \leftarrow \tau_{i-1}^{(1)} \wedge \tau_{i-1}^{(2)}$ 
6:    $t_i \leftarrow t_{i-1} + \tau_{i-1}$ 
7:   if  $\tau_{i-1}^{(1)} < \tau_{i-1}^{(2)}$  then
8:      $X_i \sim P(X_{i-1}, \cdot)$ 
9:   else
10:     $X_i \sim \mu$ 
11: end while
12: return pairs  $(t_i, X_i)$ 

```

the bouncy particle sampler in the supplementary material of [Bouchard-Côté, Vollmer and Doucet \(2018\)](#).

By conditioning on the first jump, we obtain the following representation (cf. equation (4.24) of [Moyal \(1957\)](#)):

$$\begin{aligned} Q_t f(x) &= e^{-\bar{\lambda}(x)t} f(x) + \int_0^t ds \bar{\lambda}(x) e^{-\bar{\lambda}(x)s} P^\mu[Q_{t-s} f](x) \\ &= e^{-\bar{\lambda}(x)t} f(x) + \int_0^t ds \bar{\lambda}(x) e^{-\bar{\lambda}(x)(t-s)} P^\mu[Q_s f](x). \end{aligned}$$

From this representation we can calculate the derivative,

$$\begin{aligned} \frac{dQ_t f(x)}{dt} &= -\bar{\lambda}(x) e^{-\bar{\lambda}(x)t} f(x) + \bar{\lambda}(x) P^\mu[Q_t f](x) \\ &\quad - \int_0^t ds \bar{\lambda}(x)^2 e^{-\bar{\lambda}(x)(t-s)} P^\mu[Q_s f](x). \end{aligned}$$

At $t = 0$ the definitions of $\bar{\lambda}$ and κ imply that $\frac{d}{dt} \pi Q_t f = 0$. The exchange of integration and differentiation is justified by the assumption $\pi(\bar{\lambda}^2) < \infty$. \square

In this setting, practical simulation of the Restore process is straightforward, *even when the regeneration rate is unbounded*, since the interarrival process Y is piecewise-constant. See [Algorithm 1](#) for one possible implementation.

[Algorithm 1](#) can be seen as a continuous-time variant of standard Metropolis–Hastings; at each iteration we “propose” a move according to $P(X_{i-1}, \cdot)$, which is either accepted or rejected, depending on two exponential clocks. Upon rejecting a move, rather than remaining at X_{i-1} instead we move to a new location drawn from μ .

4. Limiting properties. In this section we consider some limiting properties of the Restore process. We will not a priori assume that X has invariant distribution π , but will work in the abstract framework of [Lemma 3](#): The underlying process Y is a right process evolving on a Radon space (E, m) , we have a locally bounded measurable function $\kappa : E \rightarrow \mathbb{R}^+$, and we a probability measure μ on E . We consider the Restore process X with these dynamics. We will write $\{P_t^\mu : t \geq 0\}$ for its semigroup.

4.1. *Central limit theorem.* We first give a central limit theorem for the Restore process. Our approach here is inspired by [Hobert et al. \(2002\)](#), who considered regenerative methods for MCMC (in discrete time).

We fix a measurable function $f : E \rightarrow \mathbb{R}$.

ASSUMPTION 23 (Central limit theorem). We assume the basic conditions of Lemma 3. Furthermore we assume that X is irreducible,

$$(17) \quad \mathbb{E}_\mu[\tau_\delta^2] < \infty,$$

and that our function $f : E \rightarrow \mathbb{R}$ satisfies

$$\mathbb{E}_\mu\left[\left(\int_0^{\tau_\delta} f(X_s) ds\right)^2\right] < \infty.$$

A sufficient condition for Assumption 23 to hold is that f is a bounded function and we have simply the second moment condition (17). In turn, a sufficient condition for (17) is that Assumption 25 holds, since in that case τ_δ can be stochastically dominated by an $\text{Exp}(\kappa)$ random variable.

Under Assumption 23 we will see that a central limit theorem holds. This can be easily done since the lifetimes of the Restore process, by construction, are independent and identically distributed.

As in the construction of Restore in Section 2, set $T_0 = 0$, let (T_n) be the successive regeneration times and let $(\tau^{(i)})$ be the lifetimes. We take the initial distribution $X_0 \sim \mu$. Set for each $i = 0, 1, 2, \dots$,

$$Z_i := \int_{T_i}^{T_{i+1}} f(X_s) ds.$$

By construction the (Z_i) are independent and identically distributed, with finite first and second moments.

We can apply the strong law of large numbers to the following numerator and denominator:

$$\frac{\int_0^{T_n} f(X_s) ds}{T_n} = \frac{\sum_{i=0}^{n-1} Z_i}{\sum_{i=0}^{n-1} \tau^{(i)}} \rightarrow \frac{\mathbb{E}_\mu[\int_0^{\tau^{(0)}} f(X_s) ds]}{\mathbb{E}_\mu[\tau^{(0)}]}$$

almost surely as $n \rightarrow \infty$.

Let us write

$$\pi[f] := \frac{\mathbb{E}_\mu[\int_0^{\tau^{(0)}} f(X_s) ds]}{\mathbb{E}_\mu[\tau^{(0)}]}.$$

When the process is ergodic, this corresponds to the invariant distribution of the Restore process. It follows immediately that the random variables

$$Z_i - \tau^{(i)} \pi[f], \quad i = 0, 1, 2, \dots$$

are independent and identically distributed and have mean 0 under \mathbb{E}_μ .

Now we set, in analogue with the expression given in [Hobert et al. \(2002\)](#),

$$(18) \quad \sigma_f^2 := \frac{\mathbb{E}_\mu[(Z_0 - \tau^{(0)} \pi[f])^2]}{(\mathbb{E}_\mu[\tau^{(0)}])^2}.$$

This numerator is finite by Assumption 23.

THEOREM 24 (Central limit theorem). *We have that*

$$(19) \quad \sqrt{n} \left(\frac{\int_0^{T_n} f(X_s) ds}{T_n} - v_\pi[f] \right) \xrightarrow{d} N(0, \sigma_f^2).$$

PROOF. The left-hand side of (19) can be written

$$\sqrt{n} \left(\frac{\sum_{i=0}^{n-1} Z_i}{\sum_{i=0}^{n-1} \tau^{(i)}} - v_\pi[f] \right) = \frac{1}{\sqrt{n}} \cdot \frac{n}{\sum_{i=0}^{n-1} \tau^{(i)}} \left(\sum_{i=0}^{n-1} (Z_i - \tau^{(i)} v_\pi[f]) \right).$$

By the strong law of large numbers and the continuous mapping theorem, $n / \sum_{i=0}^{n-1} \tau^{(i)}$ converges almost surely to $(\mathbb{E}_\mu[\tau^{(0)}])^{-1}$, and in distribution also.

Hence by applying Slutsky’s lemma and the central limit theorem to the independent and identically distributed mean zero random variables $(Z_i - \tau^{(i)})$, we see that (19) holds. \square

Let us write $\bar{\tau}_n := n^{-1} \sum_{i=0}^{n-1} \tau^{(i)}$ and $\bar{f}_n := \frac{\int_0^{T_n} f(X_s) ds}{T_n}$. Similar to [Hobert et al. \(2002\)](#), our σ_f^2 can be consistently estimated by

$$\hat{\sigma}_f^2 := \frac{\sum_{i=0}^{n-1} (Z_i - \bar{f}_n \tau^{(i)})^2}{n \bar{\tau}_n^2}.$$

This is because the difference between $\hat{\sigma}_f^2$ and

$$\frac{\sum_{i=0}^{n-1} (Z_i - \tau^{(i)} v_\pi[f])^2}{n \bar{\tau}_n^2}$$

converges to zero almost surely as $n \rightarrow \infty$, and the latter is a consistent estimator for σ_f^2 .

We can use this to get an estimate of the efficiency of Restore. If we let

$$v_\pi(f) := \int (f(x) - v_\pi[f])^2 d\pi(x),$$

then we can set the effective sample size n_{eff} to be

$$n_{\text{eff}} := \frac{v_\pi(f)}{\sigma_f^2},$$

which we may be able to estimate.

We see from (18), that the denominator $(\mathbb{E}_\mu[\tau^{(0)}])^2$ will have a significant influence on the overall variance. If $\mathbb{E}_\mu[\tau^{(0)}]$ is small, the resulting variances of individual lifetimes may be unacceptably large, and as such practically speaking it is important to choose the regeneration distribution in such a way that the lifetimes are (on average) not too short. In particular, this means choosing μ which avoids regions where the regeneration rate is particularly high.

4.2. *Coupling from the past.* Under additional (fairly strong) conditions, we will have direct access to the stationary distribution of the Restore process.

ASSUMPTION 25 (Coupling from the past). There exists some $\underline{\kappa} > 0$ such that m -almost everywhere,

$$\kappa \geq \underline{\kappa} > 0.$$

We write $\|\cdot\|_\infty$ for the sup norm of a bounded function and $\|\cdot\|_1$ for the total variation norm signed measures; given a signed measure ν ,

$$\|\nu\|_1 = \sup\{|\nu(f)| : f \text{ bounded, measurable, } \|f\|_\infty \leq 1\}.$$

PROPOSITION 26 (Uniform ergodicity). *Assume the basic conditions of Lemma 3 hold, and that X is irreducible. Under Assumption 25, the Restore process X is uniformly geometrically ergodic, meaning that there exists a unique invariant distribution π such that*

$$\| \nu P_t^\mu - \pi \|_1 \leq 2e^{-t\kappa},$$

for any initial distribution ν and $t \geq 0$.

PROOF. Fix any two arbitrary initial distributions ν_1, ν_2 on E . By Assumption 25 and Poisson superposition, we can decompose the Poisson process of regeneration times as the superposition of two independent Poisson processes: a *homogeneous* Poisson process N_1 of rate $\underline{\kappa}$, and an inhomogeneous Poisson process N_2 with rate function $t \mapsto \kappa(X_t) - \underline{\kappa}$. Thus we can couple two copies of the Restore process X , with initial distributions ν_1 and ν_2 respectively, by constructing them to have N_1 in common, and the same regeneration locations. The two processes will then meet at the first arrival time of N_1 and evolve identically thereafter.

Hence by the well-known coupling inequality (see, for instance, Thorisson (2000), Section 1.5.4),

$$(20) \quad \| \nu_1 P_t^\mu - \nu_2 P_t^\mu \|_1 \leq 2e^{-t\kappa}.$$

The Markov property (i.e., the semigroup property) then shows that for any initial distribution ν , $(\nu P_t^\mu)_{t \geq 0}$ forms a Cauchy sequence in the space of probability measures equipped with the total variation norm. By completeness, there exists a limiting probability distribution π , which must also be a stationary distribution, by the Markov property and the fact that P_t^μ is a contraction in $\| \cdot \|_1$. That is, we have $\pi P_t^\mu = \pi$ for any $t \geq 0$. By irreducibility, this invariant distribution is unique. Thus taking $\nu_1 = \nu$ and $\nu_2 = \pi$ in (20) the Proposition is proven. \square

In fact under Assumption 25 we can do even better than uniform ergodicity and employ *coupling from the past* (CFTP), a technique pioneered by Propp and Wilson (1996) to obtain exact draws from the stationary distribution π . For a related approach to exact MCMC methods, see the recent approach of Jacob, O’Leary and Atchadé (2020) using couplings.

THEOREM 27 (Coupling from the past). *Under the conditions of Lemma 3 and Assumption 25, consider the Restore process X with interarrival dynamics Y , modified regeneration rate*

$$\kappa' := \kappa - \underline{\kappa} \geq 0,$$

and regeneration density μ . Suppose X is irreducible and has initial distribution

$$X_0 \sim \mu.$$

Let $T \sim \text{Exp}(\underline{\kappa})$ be independent of X . Then

$$X_T \sim \pi,$$

where π is the unique invariant distribution of the process.

PROOF. This follows from the technique of Propp and Wilson (1996). We saw in the proof of Proposition 26 that we can realise the Poisson process of regeneration times as the superposition of two independent Poisson processes: a homogeneous Poisson process N_1 of rate $\underline{\kappa}$ and an inhomogeneous Poisson process N_2 with rate $t \mapsto \kappa'(X_t)$. As in Propp and Wilson (1996), we imagine a Restore process X , initialised from some arbitrary initial distribution at time $-\infty$, run until time 0. Since we have established uniform ergodicity in Proposition 26, we know that $X_0 \sim \pi$. Let $-T$ be the most recent arrival of N_1 before time 0.

Algorithm 2 Bounded Restore Sampler: $\kappa \leq M$, with CFTP

```

1: draw run time:  $T \sim \text{Exp}(\underline{\kappa})$ 
2: initialize:  $X_0 \sim \mu, t_0 = 0, i = 0$ 
3:  $i \leftarrow i + 1$ 
4:  $t_i \leftarrow t_{i-1} + \tau_{i-1}$ , where  $\tau_{i-1} \sim \text{Exp}(M - \underline{\kappa})$ 
5: while  $t_i < T$  do
6:   simulate  $Z_i \sim \mathcal{L}(Y_{\tau_{i-1}} | Y_0 = X_{i-1})$ 
7:   with probability  $1 - (\kappa(X_i) - \underline{\kappa}) / (M - \underline{\kappa})$ 
8:      $X_i \leftarrow Z_i$ 
9:   else
10:     $X_i \sim \mu$ 
11:     $i \leftarrow i + 1$ 
12:     $t_i \leftarrow t_{i-1} + \tau_{i-1}$ , where  $\tau_{i-1} \sim \text{Exp}(M - \underline{\kappa})$ 
13: end while
14: simulate  $Z \sim \mathcal{L}(Y_{T-t_{i-1}} | Y_0 = X_{i-1})$ 
15: return  $Z$ , which is drawn exactly from  $\pi$ 

```

Regardless of the prior evolution of X , we know that $X_{-T} \sim \mu$ as $-T$ was a regeneration time. Since N_1 and N_2 are independent, X_0 then has the same law as a Restore process at time T , initialised from μ , with regeneration rate κ' .

Since the time reverse of a homogeneous Poisson process is also a homogeneous Poisson process, we can instead imagine initialising $X_0 \sim \mu$ and evolving an exponential time T into the future with modified regeneration rate κ' . \square

In the case when κ is bounded above, one implementation is given in Algorithm 2. In this case, simulation of the lifetimes τ_δ is straightforward, since it can make use of Poisson thinning; see, for instance, Devroye ((1986), Chapter 6.2).

This CFTP implementation can be seen as a continuous-time version of the multigamma coupler of Murdoch and Green (1998) or of the hybrid scheme of Murdoch ((2000), Section 3). The multigamma coupler of Murdoch and Green (1998) assumes we have a discrete-time π -invariant Markov chain whose transition kernel P satisfies $P(x, dy) = f(y|x) dy$, where $f(y|x) \geq r(y)$, for all x , for some nonnegative function r which satisfies $\rho := \int r(y) dy > 0$. Let ν_r denote the probability distribution with density (proportional to) r . Thus when simulating the chain, at each step with probability ρ , the chain will move to a point drawn from ν_r , independent of the current location. This enables a CFTP construction, the multigamma coupler; see Murdoch and Green ((1998), Section 2.1).

This uniform probability ρ is precisely what enables CFTP to be applied. It informally says that independent of location, at each discrete time step all locations are trying to couple with probability ρ to the same point, drawn from ν_r . This plays the same role as our homogeneous rate $\underline{\kappa}$, which informally states that in continuous time, at rate $\underline{\kappa}$, all locations are trying to couple to the same location, drawn from ν_μ .

A crucial difference between our approaches, however, is that our underlying dynamics Y do not themselves have to be π -invariant; in fact we will see in Section 6 an example where the local process does not possess an invariant distribution at all.

4.2.1. *Example: Classical rejection sampler.* We show that the classical rejection sampler can be seen as a special case of the CFTP implementation of the Restore process. A similar result was established for the Independence Sampler in Murdoch and Green (1998).

Let π, μ be density functions on E with respect to m . We take Y to be the trivial stochastic process on E which given its initial position Y_0 , has constant sample paths: almost surely, $Y_t = Y_0$ for all $t \geq 0$. Define the regeneration rate

$$(21) \quad \kappa(x) = C \frac{\mu(x)}{\pi(x)}, \quad x \in E,$$

for any constant $C > 0$. If we were to implement the classical rejection sampler targeting π from μ we would require the following condition:

$$(22) \quad \pi(x) \leq M\mu(x), \quad x \in E,$$

for some (finite) constant M . The classical rejection sampler targeting π from μ repeatedly draws X_n independently from μ , and accepts it with probability $\pi(X_n)/(M\mu(X_n))$, otherwise rejects it and tries again with a new $X_{n+1} \sim \mu$. The final accepted value X_n is an exact draw from π .

THEOREM 28. *Under (22), the CFTP implementation of the Restore process (Theorem 27) with constant interarrival dynamics, regeneration rate κ as in (21) and regeneration density μ is identical to classical rejection sampling targeting π from μ .*

PROOF. We see that (22) holds if and only if Assumption 25 holds with

$$\underline{\kappa} = C/M.$$

Under this condition in the CFTP implementation (Theorem 27) we run the Restore process with regeneration rate

$$\kappa' = \kappa - \underline{\kappa} = C \frac{\mu}{\pi} - \frac{C}{M}$$

for a time $T \sim \text{Exp}(C/M)$.

We can simulate this Restore process iteratively by drawing for each n , $X_n \sim \mu$. We have two competing independent exponential clocks, $T \sim \text{Exp}(C/M)$ and $T_n \sim \text{Exp}(\kappa'(X_n))$.

If $T < T_n$, all trajectories have coupled and so we terminate the algorithm and output X_n , which is an exact draw from π . By the theory of competing exponentials this occurs with probability

$$\frac{C/M}{C/M + C(\frac{\mu(X_n)}{\pi(X_n)} - \frac{1}{M})} = \frac{\pi(X_n)}{M\mu(X_n)}.$$

This is exactly the probability of acceptance for the classic rejection sampler.

If $T \geq T_n$ then we iterate again and draw $X_{n+1} \sim \mu, T_{n+1} \sim \text{Exp}(\kappa'(X_{n+1}))$. By the memoryless property of the exponential distribution we have again two independent exponential clocks as before. \square

If (22) doesn't hold, provided there is a unique invariant distribution π we can still use ergodic averages to estimate $\nu_\pi[f]$ for any bounded f . Suppose we run the Restore process with constant interarrival dynamics, regeneration rate κ as in (21) and regeneration density μ for n complete lifetimes. The corresponding ergodic average is

$$\frac{1}{T_n} \sum_{i=1}^n f(X_i)\tau^{(i)},$$

where $X_i \sim \mu$ are i.i.d., conditional on $X_i, \tau^{(i)} \sim \text{Exp}(C\mu(X_i)/\pi(X_i))$ are independent and $T_n = \sum_{i=1}^n \tau^{(i)}$. Thus the estimator of $\pi[f]$ can be seen as an importance sampling-type estimator with randomized importance weights; note $C\mathbb{E}[\tau^{(i)}|X_i] = \pi(X_i)/\mu(X_i)$.

5. Practical considerations. We consider now some practical questions related to the Restore process.

5.1. *Minimal regeneration distribution.* In this section we assume that we are given some fixed interarrival process, a positive target density π on E and a regeneration density μ on E , which are both normalized.

The most significant challenge for implementing the Restore sampler is to ensure that the regeneration rate is nonnegative; we need to find a constant C so that

$$(23) \quad \kappa(x) = \tilde{\kappa}(x) + C \frac{\mu(x)}{\pi(x)} \geq \underline{\kappa} \geq 0 \quad \text{for all } x \in E,$$

for some nonnegative constant $\underline{\kappa}$. Here $\tilde{\kappa}$ is defined in (10) for the diffusion setting and for the jump process setting is defined to be the first term on the right-hand side of (16). As shown in the proof of Theorem 16, $C = \mathbb{E}_\mu[\tau_\partial]$ can be interpreted as the average lifetime when started from μ .

One natural way to choose the regeneration density μ and constant C is to minimize the number of regeneration events. That is, we would like to choose some *minimal* regeneration distribution μ^* and constant C^* such that the regeneration rate is given by

$$(24) \quad \kappa^* := \tilde{\kappa} + C^* \frac{\mu^*}{\pi} = \tilde{\kappa} \vee \underline{\kappa}.$$

This is entirely analogous to the choice of bounce rate for the bouncy particle sampler of Bouchard-Côté, Vollmer and Doucet (2018), and of the canonical switching rate for the zig-zag in Bierkens, Fearnhead and Roberts (2019). In order to satisfy (24), the appropriate choice of density μ^* with respect to the measure m on E is

$$(25) \quad \mu^*(x) := (C^*)^{-1} [0 \vee (\underline{\kappa} - \tilde{\kappa}(x))] \pi(x),$$

where

$$C^* := \int_E [0 \vee (\underline{\kappa} - \tilde{\kappa}(x))] \pi(x) \, dm(x),$$

assuming that this quantity is finite.

PROPOSITION 29 (Minimal regeneration distribution). *Let μ^*, C^* be defined as above for some fixed $\underline{\kappa} \geq 0$, where we assume μ^* is integrable and normalized. Let μ, C be any (normalized) probability measure on E and positive constant respectively such that (23) holds. Then μ^* minorizes μ , in the sense that there exists some $\varepsilon > 0$ such that for all measurable $B \subset E$,*

$$(26) \quad \mu(B) \geq \varepsilon \mu^*(B),$$

and we have that

$$C \geq C^*.$$

PROOF. From the assumption that (23) holds, we must have that $\kappa \geq \kappa^*$ pointwise, from which it follows that for each $x \in E$,

$$C\mu(x) \geq C^*\mu^*(x),$$

which establishes (26), and by integrating both sides over E it follows that $C \geq C^*$. \square

How one can obtain samples from μ^* is in general not obvious, and is reminiscent of sampling from minorising measures as in Murdoch and Green (1998). μ^* is generally compactly supported and supported around the modes of π ; its support is contained within the

set $\{x \in E : \tilde{\kappa}(x) < \underline{\kappa}\}$, and so often simulation is possible through straightforward rejection sampling.

On the other hand, the computation of κ^* is immediate, since it does not require knowledge of C^* or μ^* but is simply a thresholded version of $\tilde{\kappa}$ as in (24).

When the interarrival process is already π -invariant, any nonnegative value of C can be used. In this setting, the recent work of Caputo and Quattropani (2019), suggests that a sensible way to tune C would be to choose it such that the average rate of regenerations matches the rate of mixing of the interarrival process. Caputo and Quattropani (2019) showed that for the similar discrete-time PageRank surfer on random (finite) graphs, the resulting mixing time depends on the interplay between the rate of mixing of the underlying walk and the regeneration probability.

5.2. *Truncated regeneration rate.* We consider now the diffusion case, as in Section 3.1. In this case κ is typically unbounded, and the simulation of the lifetimes $\tau^{(i)}$ is not straightforward. In some cases using layered processes it is still possible to simulate τ_∂ exactly, as with the techniques of Pollock et al. (2020). These are technically demanding, so in this section we consider the alternative of *truncating* the regeneration rate. Namely, we fix some upper bound M , and work with the truncated regeneration rate

$$\kappa_M := \kappa \wedge M.$$

This will introduce some approximation error, a discrepancy between the invariant distribution and π , but we will show how this error may be explicitly quantified.

In order to prove our result we will need to assume the following.

We assume that the interarrival process Y is a diffusion on \mathbb{R}^d satisfying BASSA, and that κ is continuous. We also assume that Assumption 25 holds, namely that we have a lower bound

$$\kappa \geq \underline{\kappa} > 0.$$

Recall that under Assumption 25, τ_∂ can be stochastically dominated by an exponential random variable with rate $\underline{\kappa}$, and hence all moments of τ_∂ are finite.

In order to avoid pathologies we assume that

$$(27) \quad M > \inf_{x \in E} \kappa(x).$$

We consider now the Restore process X with interarrival process Y , regeneration density μ and truncated regeneration rate κ_M , for some given truncation level M satisfying (27).

Throughout this section we will be concerned only with the behavior of the Restore process before the first regeneration event. As the regeneration distribution μ will not play a significant role we will consider the local process Y , without regenerations, and explicitly augment it with a first regeneration time. We will simply write \mathbb{E}_x for the law of the local process Y started from x , and consider the first arrival time τ_∂ to be a random variable defined by (3).

Let us write κ_M^c for the excess regeneration rate over level M , that is,

$$\kappa_M^c := \kappa - \kappa_M.$$

Then by Poisson superposition, we can write

$$(28) \quad \tau_\partial = \tau_M \wedge \tau_M^c,$$

where $\tau_\partial, \tau_M, \tau_M^c$ are the first arrival times of inhomogeneous Poisson process with rate functions $t \mapsto \kappa(Y_t), t \mapsto \kappa_M(Y_t)$ and $t \mapsto \kappa_M^c(Y_t)$ respectively, where these latter two Poisson processes are independent conditional on the path $t \mapsto Y_t$.

In particular, τ_M and τ_M^e can be written as

$$(29) \quad \tau_M = \inf \left\{ t \geq 0 : \int_0^t \kappa_M(Y_s) \, ds \geq \xi_1 \right\},$$

$$(30) \quad \tau_M^e = \inf \left\{ t \geq 0 : \int_0^t \kappa_M^e(Y_s) \, ds \geq \xi_2 \right\},$$

where $\xi_1, \xi_2 \sim \text{Exp}(1)$ are independent of each other and of the underlying process Y .

Since we are assuming Assumption 25 holds, by the arguments of Section 4.2 it follows that the Restore process with regeneration rate κ has a unique invariant distribution π , and from Section 4.1 the action of π on a test function f can be written as

$$\nu_\pi[f] = \frac{\mathbb{E}_\mu[\int_0^{\tau_\partial} f(Y_s) \, ds]}{\mathbb{E}_\mu[\tau_\partial]},$$

where here Y is the local process without regenerations and τ_∂ is defined as in (3).

Similarly, the Restore process with truncated regeneration rate κ_M is still uniformly ergodic and possesses a unique invariant distribution π_M .

Our goal now is to bound the total variation distance

$$\|\pi_M - \pi\|_1,$$

as a function of M .

THEOREM 30. *We have the following bound on the error:*

$$\|\pi_M - \pi\|_1 \leq \frac{4 \int_0^\infty \mathbb{P}_\mu(\tau_M^e \leq t) \exp(-t\kappa) \, dt}{\mathbb{E}_\mu[\tau_\partial]}.$$

PROOF. See Appendix A.5. \square

REMARK 31. To use this bound we need to further bound

$$\mathbb{P}_x(\tau_M^e \leq t).$$

Intuitively, if κ_M is a reasonable approximation for κ , then κ_M^e is low, and hence τ_M^e tends to be large, and so this bound is tighter.

PROPOSITION 32. *Fix a regeneration distribution μ . We have that*

$$(31) \quad \int_0^\infty dt \mathbb{P}_\mu(\tau_M^e \leq t) e^{-t\kappa} \rightarrow 0 \quad \text{as } M \rightarrow \infty.$$

Thus by Theorem 30 as $M \rightarrow \infty$,

$$\|\pi_M - \pi\|_1 \rightarrow 0.$$

PROOF. The event $\{\tau_M^e \leq t\}$ is contained in the event $\{\sup_{s \leq t} \kappa(Y_s) \geq M\}$. Thus, for any fixed x

$$\begin{aligned} \lim_{M \rightarrow \infty} \mathbb{P}_x(\tau_M^e \leq t) &= \mathbb{P}_x \left(\bigcap_{M=1}^\infty \{\tau_M^e \leq t\} \right) \\ &\leq \mathbb{P}_x \left(\sup_{s \leq t} \kappa(Y_s) = \infty \right) \\ &\leq \mathbb{P}_x \left(\sup_{s \leq t} \|Y_s\| = \infty \right) \quad \text{since } \kappa \text{ is locally bounded} \\ &= 0 \quad \text{since } Y \text{ is nonexplosive.} \end{aligned}$$

By the dominated convergence theorem it follows that

$$\lim_{M \rightarrow \infty} \int_0^\infty dt e^{-t\kappa} \int_E d\mu(x) \mathbb{P}_x(\tau_M^c \leq t) = 0,$$

which is precisely (31). \square

In order for Theorem 30 to be of practical use, we will further need bounds on

$$(32) \quad \mathbb{P}_x(\tau_M^c \leq t),$$

which will vary given the particular situation; given the choice of the underlying diffusion Y , target π and regeneration density μ .

The rate at which the probabilities (32) decay as a function of M will crucially depend on the rate at which the regeneration rate κ grows. Thus we define the following,

$$L(M) := \sup\{\ell > 0 : \sup\{\kappa(x) : x \in [-\ell, \ell]^d\} \leq M\},$$

which for a given truncation level M defines the largest hypercube on which no truncation occurs.

The rate at which $L(M)$ grows as $M \rightarrow \infty$ will crucially dictate the rate at which the error decays. Then let

$$H(M) := [-L(M), L(M)]^d \subset \mathbb{R}^d,$$

and let

$$T_M := \inf\{t \geq 0 : Y_t \in \mathbb{R}^d \setminus H(M)\}$$

be the first hitting time of the diffusion Y (without regenerations) of the complement of $H(M)$. Clearly we must have

$$T_M \leq \tau_M^c.$$

Thus it follows that

$$\int_0^\infty \mathbb{P}_x(\tau_M^c \leq t) e^{-\kappa t} dt \leq \int_0^\infty \mathbb{P}_x(T_M \leq t) e^{-\kappa t} dt.$$

To proceed from here we require knowledge of the distribution of the hitting times T_M for the underlying diffusion Y . At this point we will *specialize to the case of Brownian motion*; however, a similar analysis can be performed in any situation where we have analogous bounds on the hitting times.

By the reflection principle for one-dimensional Brownian motion we know that for any $a > 0$,

$$\begin{aligned} \mathbb{P}\left(\sup_{0 \leq s \leq t} |B_s| > a\right) &\leq 2\mathbb{P}\left(\sup_{0 \leq s \leq t} B_s > a\right) = 4\mathbb{P}(B_t > a) \\ &= 4\left(1 - \Phi\left(\frac{a}{\sqrt{t}}\right)\right). \end{aligned}$$

Here Φ denotes the standard univariate normal cumulative distribution function. For a multi-dimensional standard Brownian motion, it follows that

$$\mathbb{P}_0(T_M \leq t) \leq 4d\left(1 - \Phi\left(\frac{L(M)}{\sqrt{t}}\right)\right).$$

This is because leaving a hypercube is the same as having some component leaving the interval $[-L(M), L(M)]$.

We now make use of the well-known bound for the standard normal cumulative distribution function: for each $\lambda > 0$,

$$1 - \Phi(\lambda) < \frac{1}{\sqrt{2\pi}\lambda} e^{-\lambda^2/2}.$$

This leads to the bound

$$\int_0^\infty \mathbb{P}_0(T_M \leq t) e^{-\kappa t} dt \leq \frac{4d}{\sqrt{2\pi}} \int_0^\infty \frac{\sqrt{t}}{L(M)} e^{-L(M)^2/(2t)} e^{-\kappa t} dt.$$

This integral can be evaluated analytically,¹ to obtain

$$\begin{aligned} \int_0^\infty \mathbb{P}_0(T_M \leq t) e^{-\kappa t} dt &\leq \frac{4d}{\sqrt{2\pi}} \frac{\sqrt{\pi}}{2\kappa} \left(\sqrt{2} + \frac{1}{L(M)\sqrt{\kappa}} \right) e^{-\sqrt{2\kappa}L(M)} \\ &= \frac{2d}{\kappa} \left(1 + \frac{1}{L(M)\sqrt{2\kappa}} \right) e^{-\sqrt{2\kappa}L(M)}. \end{aligned}$$

So for large values of M we have a bound that decays like

$$e^{-\sqrt{2\kappa}L(M)}.$$

This can be used to give practical suggestions of how large to choose M in order to balance the bias and variance of the algorithm’s output.

Suppose we are able to obtain n i.i.d. draws $X_1, \dots, X_n \sim \pi^M$, say by running the CFTP algorithm a total of n times. For a bounded test function f , we estimate $v_\pi[f]$ by

$$\sum_{i=1}^n \frac{f(X_i)}{n}.$$

We estimate the error roughly as

$$\begin{aligned} \left| \sum_{i=1}^n \frac{f(X_i)}{n} - v_\pi[f] \right| &\leq \underbrace{\left| \sum_{i=1}^n \frac{f(X_i)}{n} - \pi^M(f) \right|}_{\sim \frac{1}{\sqrt{n}}} + \underbrace{|\pi^M(f) - v_\pi[f]|}_{\leq \|f\|_\infty \|\pi^M - \pi\|_{\text{TV}}} \\ &\approx O\left(\frac{1}{\sqrt{n}}\right) + \exp(-\sqrt{2\kappa}L(M)). \end{aligned}$$

In order to balance these two terms, it is advisable to choose n and M such that

$$\begin{aligned} \frac{1}{\sqrt{n}} &\sim \exp(-\sqrt{2\kappa}L(M)) \\ \Rightarrow \frac{\log n}{2\sqrt{2\kappa}} &\sim L(M). \end{aligned}$$

So this gives some indication of how to choose M , given n . This will achieve an error of order roughly $O(n^{-1/2})$. The computational cost in n will be roughly $O(n \log n)$.

6. Examples. In this section we give some univariate examples which highlight key aspects of our Restore methodology. A thorough investigation of the computational properties of Restore is an important and challenging task, which is outside the scope of this present work and will be the topic of future research.

¹[https://www.wolframalpha.com/input/?i=int_0%5Einfty+%5Csqrt+\(t\)+exp\(-a%2F\(2t\)\)+exp\(-t\)dt](https://www.wolframalpha.com/input/?i=int_0%5Einfty+%5Csqrt+(t)+exp(-a%2F(2t))+exp(-t)dt)

6.1. *Cauchy posterior.* We first give an example where π has heavy tails and is multi-modal, where we can apply coupling from the past.

This example is based on Example 3.1 of Murdoch (2000). We take

$$(33) \quad \bar{\pi}(x) \propto \prod_{i=1}^n \frac{1}{1 + (y_i - x)^2},$$

for some observations $(y_1, \dots, y_n) \in \mathbb{R}^n$, with respect to Lebesgue measure on \mathbb{R} . (We use the notation $\bar{\pi}$, since in the notation of Section 3.1, the symbol π is reserved for the target density with respect to the measure Γ .)

This can be thought of as the posterior distribution for i.i.d. $\text{Cauchy}(x)$ data, with an improper uniform prior on \mathbb{R} for x . In Example 3.1 of Murdoch (2000), the author considers a very similar target with lighter tails. We will take the same data as Murdoch (2000), namely $n = 3$ and observations $(1.3, -11.6, 4.4)$. The resulting posterior is plotted in red in Figure 1. Our sampling approach here is similar to that of Murdoch (2000); we are also combining local and global dynamics, but we will choose diffusive local dynamics which rapidly enter the tails.

As such, for our underlying process, we will take the following diffusion: an *unstable* Ornstein–Uhlenbeck process, described by the SDE

$$(34) \quad dY_t = Y_t dt + dB_t,$$

where B is a standard univariate Brownian motion. We showed in Section 3.1.1 that this diffusion satisfies the BASSA conditions (Assumption 7). This diffusion, like a stable Ornstein–Uhlenbeck process, is also a Gaussian process with known finite-dimensional distributions, and so can be simulated easily without error.

For the regeneration distribution we will take the minimal regeneration distribution μ^* from Section 5.1, with $\kappa = 4$. This distribution is compactly supported, and samples can be

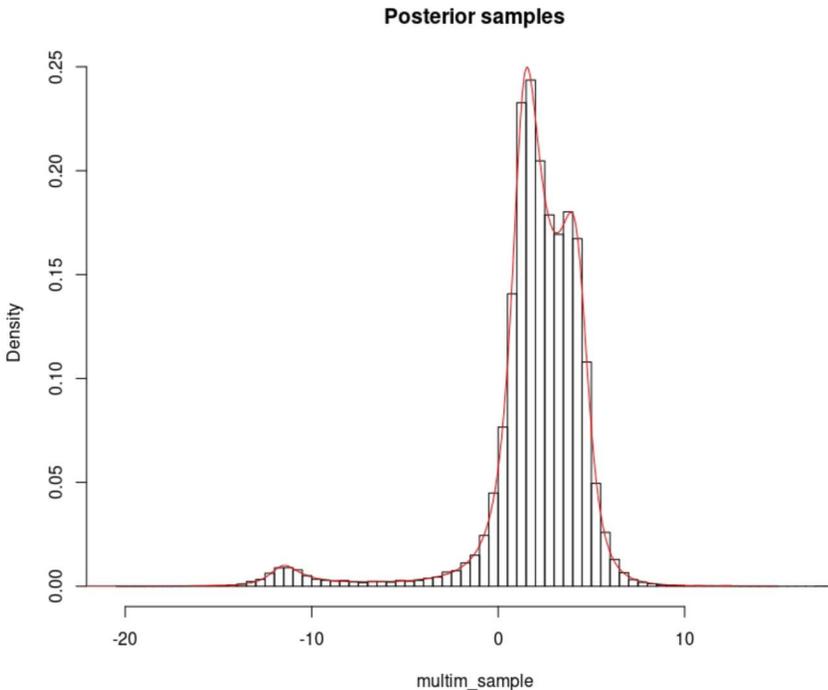


FIG. 1. The heavy-tailed multi-modal target distribution π of (33) (red), and 30,000 samples obtained from the CFTP implementation. These are i.i.d. draws from π .

efficiently obtained through rejection sampling from a uniform distribution. In this setting the regeneration rate is uniformly bounded from above, and so we can directly make use of Poisson thinning, as in Algorithm 2. The various assumptions as in Section 3.1 are easily verified to hold in this setting.

Thus we are able to apply the CFTP implementation (Section 4.2) to obtain independent and identically distributed draws from π . A histogram consisting of 30,000 draws from the CFTP implementation are plotted in Figure 1. These were obtained by running the CFTP algorithm 30,000 times independently.

6.2. Jump Restore example. We turn now to an example of jump process Restore (Section 3.2), where we use Restore to introduce rejection-free moves into an existing sampler.

A situation where Assumption 19 is easily checked is when P corresponds to a Markov chain that is already π -invariant, for instance the kernel of an appropriate MCMC algorithm targeting π . In this case we can easily embed P into continuous time without changing the asymptotic dynamics, just by taking constant holding rates $\lambda \equiv 1$. In this case the regeneration rate reduces to

$$\kappa(x) = C \frac{\mu(x)}{\pi(x)}, \quad x \in E,$$

and we see that any choice of $C > 0$ will ensure nonnegativity of κ . This gives a recipe to introduce rejection-free moves to a discrete sampler in continuous time.

Consider the following example, in one dimension for ease of visualisation. Writing $\phi(\cdot; \nu, \sigma^2)$ for the univariate Gaussian density with mean $\nu \in \mathbb{R}$ and variance $\sigma^2 > 0$, take as the target π on \mathbb{R} :

$$\pi(x) = 0.1 \phi(x; -22, 3^2) + 0.3 \phi(x; -1, 0.2^2) + 0.6 \phi(x; 15, 1^2), \quad x \in \mathbb{R}.$$

For the regeneration density μ , we take

$$\mu(x) = \frac{1}{3} (\phi(x; -29, 0.3^2) + \phi(x; 3, 1^2) + \phi(x; 10, 1^2)), \quad x \in \mathbb{R}.$$

We take the underlying process Y to be random walk metropolis with variance 1 embedded in continuous time, with constant holding rate 1. We took the constant $C = 1$ in the regeneration rate.

We have plotted a histogram after 300,000 steps of the jump chain (taking into account holding times) in Figure 2 and in Figure 3 we have plotted the continuous-time trajectory of the first 50,000 jump steps of this run.

7. Conclusions. In this work we have introduced and studied the *Restore process*, which is obtained by enriching an existing *local* continuous-time Markov process with *global* regenerative dynamics. We have focused particularly on how it could be applied within Monte Carlo methodology to sample from a prescribed target density of interest, π . Surprisingly, the Restore process enables us to combine continuous-time local and global dynamics—*neither of which is π -invariant*—and by means of an inhomogeneous Poisson process compensate these dynamics to ensure the process is π -stationary. The resulting sampler is simple to implement (in many settings it is no more complex than a vanilla Metropolis–Hastings sampler), and we readily establish a central limit theorem. Although the use of an inhomogeneous Poisson process has some natural affinity with recent Monte Carlo developments (such as piecewise-deterministic MCMC methods, Bouchard-Côté, Vollmer and Doucet (2018), Bierkens, Fearnhead and Roberts (2019) and quasi-stationary Monte Carlo methods, Pollock et al. (2020), Wang et al. (2019)), the additional regenerative behavior provides significant

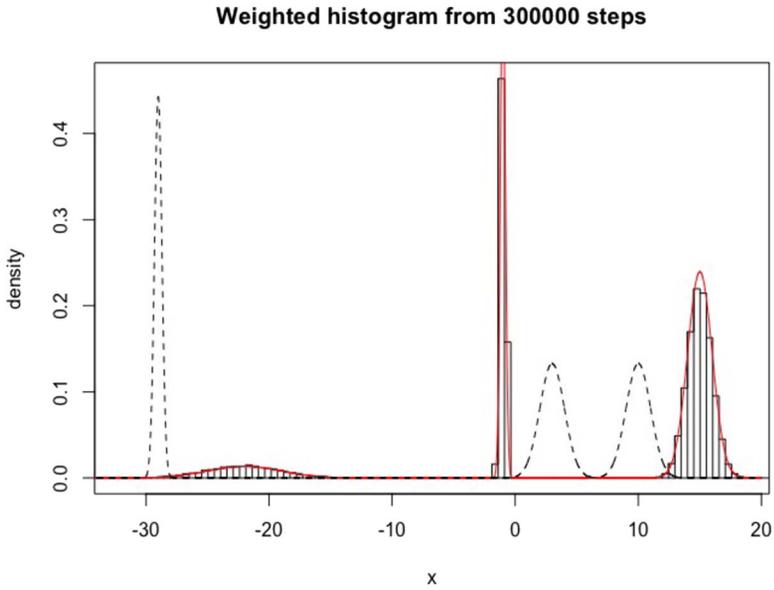


FIG. 2. The target density π (red) and regeneration density μ (dashed) for the jump Restore example, along with a weighted histogram of the Restore run, taking into account holding times.

benefits, both theoretically and practically. In particular, we show that the regenerative behavior allows us, with verifiable conditions, to demonstrate uniform ergodicity, avoid traditional MCMC problems such as burn-in, and even construct independent exact draws from π by a coupling-from-the-past scheme. On the more practical side, we have discussed some of the natural approximations a practitioner may make in implementing continuous-time samplers for use within the Restore process, showing that the global regenerative dynamics allow us to readily analyze such approximations, and to study and understand their effect.

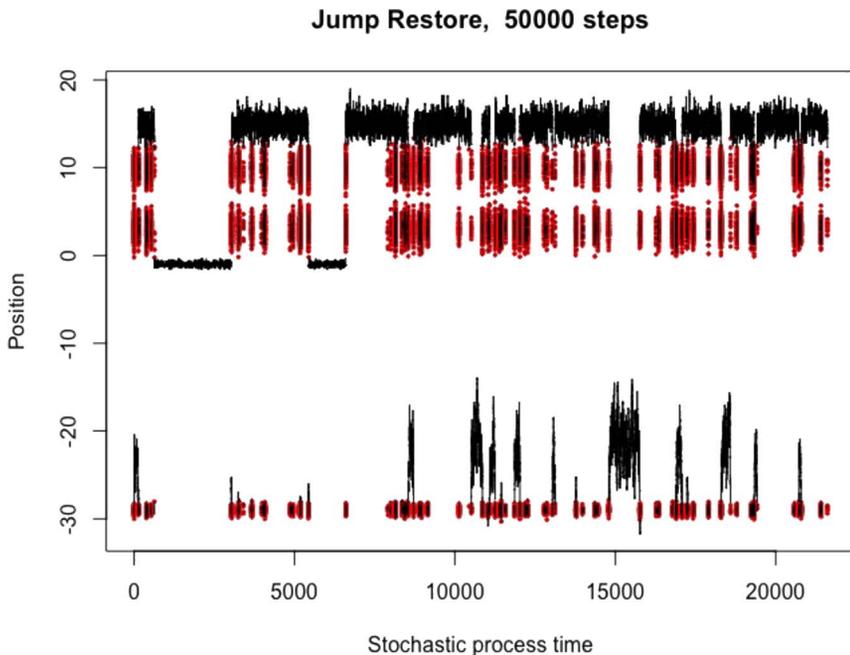


FIG. 3. The continuous-time trajectory of the first 50,000 steps of the jump Restore sampler. The proportion of regeneration moves was roughly 0.496. The red points are the regenerations.

This work has been primarily focused on introducing and establishing fundamental results for the Restore process, and there is considerable scope for further development. From a practical perspective, it is of interest to consider its use in different settings, with new classes of local and global dynamics. For instance, in the setting of continuous-time local dynamics one could consider piecewise-deterministic Markov processes, or even ODE flows, which will have properties particularly suited to certain problems. Or the construction of the global dynamics—the regeneration density—might utilize *other* inference about the target density π , obtained, say, by a simpler but approximate scheme. Thus in effect one could use Restore to remove the error from existing statistical approaches. Indeed, the flexibility offered by Restore, in combining continuous-time local and global dynamics which are not independently π -invariant, suggests that we are still far from espying the limits to which this general framework might be extended. For instance, it may be possible to use the framework to compensate multiple local and global dynamics, or to have global dynamics which are themselves adaptive to the accrued information of the process. Other prospective applications of the Restore process include exploiting the regenerative structure of the process for use within parallel computing architectures, embedding Restore within other Monte Carlo methodologies which require independent exact draws from the target distribution, something that Restore (unlike MCMC) can provide.

From a theoretical perspective, our understanding of the Restore process is nowhere near as complete as we should like, beginning with the lack of a single unified proof of invariance of the target distribution, and proceeding to the need for an appropriate definition of ‘efficiency’ or ‘optimality’ for choosing appropriate dynamics. Any reckoning with these notions would have to take account of the temporally varying computational cost of simulating the process with multiple dynamics, which is beyond the scope of this paper. Insights into issues such as these could be particularly useful in the design of appropriate diagnostics for the process, for instance, in situations where the verifiable CFTP conditions do not hold.

APPENDIX: PROOFS

A.1. Proof of Lemma 3. The techniques of Chapter 61 of Sharpe (1988) allow us to identify the process Y killed at time (3) with the (sub-)process generated by the decreasing multiplicative functional

$$m_t := \exp\left(-\int_0^t \kappa(Y_s) ds\right), \quad t \geq 0.$$

Since Y is right-continuous and κ is locally bounded, this defines a *right* multiplicative functional. Then by Theorem 61.5 of Sharpe (1988), concerning processes which are generated by such right multiplicative functionals, we can conclude that our killed process Y with lifetime (3) is a right process.

The resulting Restore process X , given in (4), is formed by concatenating independent copies of such killed processes with initial distribution μ . Exercise 14.17 of Sharpe (1988) shows that the infinite concatenation of a series of independent and identically distributed right processes is yet another right processes. Hence our Restore process X is indeed a right process.

Since κ is locally bounded and Y is right-continuous, it follows that $\mathbb{E}_\mu[\tau_\partial] > 0$. The final statement then follows from the fact that the lifetimes (apart from possibly the first) are independent and identically distributed.

A.2. Proof of Proposition 12. We have seen in the Proof of Lemma 3 that the killed process can be identified with a subprocess generated by a multiplicative functional. Hence we will seek to utilize Theorem 2.5 of Demuth and van Casteren (2000), which can be applied

to such processes. From our Assumptions 8 and 11, it follows that κ is nonnegative and continuous. In particular it is bounded on compact sets. Thus the potential $V = \kappa$ belongs to the so-called *local Kato class*. We are assuming that our underlying process Y satisfies BASSA (Assumption 7), so we have satisfied the conditions of Theorem 2.5 of Demuth and van Casteren (2000).

The conclusions of the Proposition are precisely the conclusions of Theorem 2.5 of Demuth and van Casteren (2000), restated in our present setting.

A.3. Proof of Lemma 15. Recall that we write $W^{2,1}(\mathbb{R}^d)$ for the Sobolev space of measurable functions on \mathbb{R}^d whose first and second derivatives are integrable with respect to Lebesgue measure on \mathbb{R}^d , equipped with the corresponding Sobolev norm. Precise definitions can be found in Adams (1975).

First note that the fact that $\bar{\pi} \in W^{2,1}(\mathbb{R}^d)$ along with the integral assumptions imply that

$$\int_{\mathbb{R}^d} \pi(x)\kappa(x) \, d\Gamma(x) < \infty,$$

since we can write

$$\pi\kappa\gamma = \frac{1}{2} \Delta \bar{\pi} - \nabla A \cdot \nabla \bar{\pi} - \Delta A \bar{\pi} + C\mu\gamma.$$

Now, as in the proof of Theorem 3.18 of Adams (1975), let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a mollifier, satisfying properties (i), (ii) and (iii) described therein with $m = 2$. Taking the square if necessary, we can assume that f is nonnegative. In particular, f and its derivatives up to order 2 are bounded pointwise in absolute value by a constant M . We can now define similarly for each $n \in \mathbb{N}$, $\pi_n := f_n\pi$, where $f_n(x) := f(x/n)$, $x \in \mathbb{R}^d$.

Since π is smooth, the π_n are a sequence of smooth, compactly supported functions with the following properties: $\bar{\pi}_n := \gamma \cdot \pi_n$ converges to $\bar{\pi} = \gamma \cdot \pi$ pointwise and in $W^{2,1}(\mathbb{R}^d)$ (as in the proof of Theorem 3.18, Adams (1975)), and we have $\pi_n \leq M\pi$ pointwise, uniformly over n . This implies, in particular, that π_n converges to π in $\mathcal{L}^1(\Gamma)$.

Using the relation $\nabla A = \frac{\nabla\gamma}{2\gamma}$ we can relate the action of L^0 on Γ -densities to its action on Lebesgue densities:

$$(35) \quad -\gamma L^0 \pi_n = \frac{1}{2} \Delta \pi_n + \nabla A \cdot \nabla \pi_n = \frac{1}{2} \Delta \bar{\pi}_n - \nabla A \cdot \nabla \bar{\pi}_n - (\Delta A) \bar{\pi}_n.$$

This will allow us to show that $-\gamma L^0 \pi_n$ converges in $\mathcal{L}^1(\mathbb{R}^d)$ — \mathbb{R}^d equipped with Lebesgue measure—to

$$\frac{1}{2} \Delta \bar{\pi} - \nabla A \cdot \nabla \bar{\pi} - \Delta A \bar{\pi}.$$

We consider the three terms on the right-hand side individually. Convergence of the first term is immediate since $\bar{\pi}_n$ converges to $\bar{\pi}$ in $W^{2,1}(\mathbb{R})$. Convergence of the third term follows since $\bar{\pi}$ converges to $\bar{\pi}$ pointwise, with $\pi_n \leq M\pi$, so we can make use of the dominated convergence theorem.

It remains to demonstrate the convergence of $\nabla A \cdot \nabla \bar{\pi}_n$. We have

$$(36) \quad \nabla A \cdot \nabla \bar{\pi}_n = f_n \nabla A \cdot \nabla \bar{\pi} + \bar{\pi} \nabla A \cdot \nabla f_n.$$

The first term on the right-hand side converges in $\mathcal{L}^1(\mathbb{R}^d)$ to $\nabla A \cdot \nabla \bar{\pi}$ straightforwardly, by dominated convergence, as the mollifiers are uniformly bounded by M . For the second term of (36), first note that $\nabla f_n(x) = n^{-1} \nabla f(x/n)$, which is bounded (in each component) by M/n , and the support of ∇f_n is by construction within the set $B_n := \{y \in \mathbb{R}^d : n \leq |y| \leq 2n\}$. Thus we have

$$|\bar{\pi}(x) \nabla A \cdot \nabla f_n(x)| \leq \bar{\pi}(x) K |x| M n^{-1} 1_{B_n}(x) \leq 2KM\bar{\pi}(x).$$

Here we used the bound on the drift $|\nabla A(x)| \leq K|x|$. Thus we can apply the dominated convergence theorem once more to establish that $\bar{\pi} \nabla A \cdot \nabla f_n$ is converging in $\mathcal{L}^1(\mathbb{R}^d)$ to the zero function.

It follows (by reversing the application of (35)) that $-L^0 \pi_n$ converges to

$$\frac{1}{2} \Delta \bar{\pi} - \nabla A \cdot \nabla \bar{\pi} - (\Delta A) \bar{\pi} = \frac{1}{2} \Delta \pi + \nabla A \cdot \nabla \pi$$

in $\mathcal{L}^1(\Gamma)$. Since L^0 is a closed operator, and since the sequence of smooth compactly supported functions π_n belongs to $\mathcal{D}(L_1^0)$ for each n , we have thus established that $\pi \in \mathcal{D}(L_1^0)$.

Finally, since $\int \pi \kappa \, d\Gamma < \infty$, and we have that $\pi_n \leq M\pi$, we have that $\int \pi_n \kappa \, d\Gamma \rightarrow \int \pi \kappa \, d\Gamma$.

Thus

$$L^\kappa \pi_n \rightarrow L^0 \pi + \kappa \pi = C\mu$$

in $\mathcal{L}^1(\Gamma)$. This shows that $\pi \in \mathcal{D}(L_1^\kappa)$ and $L_1^\kappa \pi = C\mu$, concluding the proof of Lemma 15.

A.4. Proof of Theorem 16. We want to prove that π is an invariant distribution for the Restore process X with interarrival dynamics Y , regeneration rate κ as defined in (9) with regeneration density μ . We are in the setting $(E, m) = (\mathbb{R}^d, \Gamma)$.

We know that the Restore process X , formed by concatenating copies of the killed process, is a strong Markov process; see Lemma 3 and its proof. Let $\{P_t^\mu : t \geq 0\}$ denotes its semigroup.

Our goal is to show

$$t \mapsto \pi P_t^\mu f := \int d\Gamma(x) \pi(x) [P_t^\mu f](x) = \mathbb{E}_\pi[f(X_t)]$$

is constant in t . By time homogeneity it suffices to show that the time-derivative is 0 at $t = 0$. This is the same method used to prove π -invariance of the bouncy particle sampler in the supplementary material of [Bouchard-Côté, Vollmer and Doucet \(2018\)](#).

The Restore process naturally exhibits renewal behavior, since the individual lifetimes are independent and identically distributed. So we will seek a renewal-type representation of the semigroup P_t^μ by conditioning on the first arrival τ_∂ . Since κ is locally bounded, τ_∂ is absolutely continuous on \mathbb{R}^+ , hence will possess a density with respect to Lebesgue measure on \mathbb{R}^+ .

Since κ is nonnegative, the semigroup $\exp(-tL^\kappa)$ can also be expressed as

$$[\exp(-tL^\kappa)f](x) = \mathbb{E}_x[f(Y_t)1\{\tau_\partial > t\}],$$

where τ_∂ is defined as in (3) for each f where the integral is well defined. Note that we have

$$(37) \quad -L^\kappa \pi = -L^0 \pi - \kappa \pi = -C\mu.$$

This equation holds formally, where we view L^κ and L^0 as formal differential operators, and as a statement about the $\mathcal{L}^1(\Gamma)$ generator by Assumption 14. Since we additionally assume that in Assumption 14 that both π and μ are in $\mathcal{L}^2(\Gamma)$ it follows that (37) also holds for the $\mathcal{L}^2(\Gamma)$ generator as well.

Consider

$$\begin{aligned} \mathbb{P}_\pi(\tau_\partial > t) &= \int d\Gamma(x) \pi(x) \int d\Gamma(y) p^\kappa(t, x, y) \\ &= \int d\Gamma(y) \int d\Gamma(x) \pi(x) p^\kappa(t, y, x) \\ &= \int d\Gamma(y) [\exp(-tL^\kappa)\pi](y), \end{aligned}$$

where the second line relies on Tonelli’s theorem to exchange the order of integration, and uses the symmetry of p^κ to replace $p^\kappa(t, x, y)$ by $p^\kappa(t, y, x)$. The final integral is well defined since $\pi \in \mathcal{L}^1(\Gamma)$ and the semigroup $\exp(-tL^\kappa)$ maps $\mathcal{L}^1(\Gamma)$ to itself, by Proposition 12. Thus by strong continuity and the fact that $\pi \in \mathcal{D}(L_1^\kappa)$ (Assumption 14) we can differentiate this expression to find

$$\begin{aligned} \frac{d}{dt} \mathbb{P}_\pi(\tau_\partial > t) \Big|_{t=s} &= \int d\Gamma(y) [\exp(-sL^\kappa)(-L^\kappa \pi)](y) \\ &= - \int d\Gamma(y) [\exp(-sL^\kappa)(C\mu)](y) \\ &= -C \int d\Gamma(x) \mu(x) [\exp(-sL^\kappa)1](x). \end{aligned}$$

The second line applies (37) again, while the final equality relies once more on Assumption 14 and symmetry of the semigroup.

This shows that the density on \mathbb{R}^+ with respect to Lebesgue measure of the first arrival time under \mathbb{P}_π is given by

$$h(s) = C \int d\Gamma(x) \mu(x) [\exp(-sL^\kappa)1](x) = C \mathbb{P}_\mu(\tau_\partial > s), \quad s \geq 0,$$

and that $C = 1/\mathbb{E}_\mu[\tau_\partial]$. This allows us to represent the semigroup of the Restore process started in π as

$$\begin{aligned} \pi P_t^\mu f &= \int_0^t C \mathbb{P}_\mu(\tau_\partial > s) \mu P_{t-s}^\mu f \, ds + \pi \exp(-tL^\kappa) f \\ &= C \int_0^t \mathbb{P}_\mu(\tau_\partial > t-s) \mu P_s^\mu f \, ds + \pi \exp(-tL^\kappa) f. \end{aligned}$$

Our goal is to differentiate this expression with respect to t , and to show that the derivative at $t = 0$ is zero.

Consider any bounded f in $\mathcal{D}(L_2^\kappa)$. From the representation above we can see that $t \mapsto \pi P_t^\mu f$ is a continuous function. Starting from

$$\begin{aligned} t \mapsto \mathbb{P}_\mu(\tau_\partial > t) &= \int d\Gamma(x) \mu(x) \int p^\kappa(t, x, y) d\Gamma(y) \\ &= \mathbb{E}_\mu^0 \left[\exp\left(-\int_0^t \kappa(Y_s) ds\right) \right], \end{aligned}$$

our technical assumption (12) allow us to differentiate under the integral sign to obtain

$$g(s) := -\frac{d}{dt} \mathbb{P}_\mu(\tau_\partial > t) \Big|_{t=s} = \mathbb{E}_\mu^0 \left[\kappa(Y_s) \exp\left(-\int_0^s \kappa(Y_u) du\right) \right]$$

for each $s \in [0, 1]$. g is a continuous function, and will be uniformly bounded over $s \in [0, 1]$.

Conditioning, as above, on the first regeneration time, we then have

$$\mu P_t^\mu f = \int_0^t g(s) \mu P_{t-s}^\mu f \, ds + \mu \exp(-tL^\kappa) f,$$

showing that $t \mapsto \mu P_t^\mu f$ is also a continuous function. By Leibniz’s rule:

$$\frac{d}{dt} \mu P_t^\mu f \Big|_{t=s} = C \int_0^s g(s-u) \mu P_u^\mu f \, du + C \mu(1) \mu P_s^\mu f + \mu \exp(-sL^\kappa) (-L^\kappa f).$$

Taking $t = 0$, we find

$$\frac{d}{dt} \mu P_t^\mu f \Big|_{t=0} = C \mu(f) + \mu (-L^\kappa f).$$

Since we chose $f \in \mathcal{D}(L_2^\kappa)$, and since we assumed $\pi \in \mathcal{D}(L_2^0) \cap \mathcal{D}(L_2^\kappa)$ (Assumption 8), this final expression is equal to

$$C\mu(f) + \pi(-L^\kappa f) = \int d\Gamma(x) f(x)(C\mu(x) + (-L^0\pi)(x) - \kappa(x)\pi(x)).$$

This will equal 0 for any such f if

$$\kappa(x) = \frac{-L^0\pi}{\pi}(x) + C\frac{\mu(x)}{\pi(x)}, \quad x \in \mathbb{R}^d,$$

which is exactly our (9). This concludes the proof of Theorem 16.

A.5. Proof of Theorem 30. Recall the expression for the invariant distribution

$$\nu_\pi[f] = \frac{\mathbb{E}_\mu[\int_0^{\tau_\partial} f(Y_s) ds]}{\mathbb{E}_\mu[\tau_\partial]}.$$

We can rewrite this by exchanging the order of integration. Consider the resolvent operator, which maps measurable functions to measurable functions,

$$\begin{aligned} \mathcal{R}f(x) &:= \int_0^\infty dt \mathbb{E}_x[f(Y_t)1\{\tau_\partial > t\}] \\ &= \int_0^\infty dt \mathbb{E}_x[f(Y_t)1\{\tau_M > t\}1\{\tau_M^e > t\}], \end{aligned}$$

where the second equality holds by (28). Note that given a bounded measurable function f , $\mathcal{R}f$ is also a bounded measurable function, since we can bound

$$|\mathcal{R}f(x)| \leq \|f\|_\infty \mathbb{E}_x[\tau_\partial] \leq \|f\|_\infty \underline{\kappa}^{-1}.$$

Thus by Fubini’s theorem, we can write

$$\pi = \frac{\mu\mathcal{R}}{\mu\mathcal{R}1},$$

in analogue with expressions given in Wang, Roberts and Steinsaltz (2020) and Benaim, Cloez and Panloup (2018).

The invariant distribution π_M of the process with truncated rate can be represented in a similar way. Write

$$\pi_M = \frac{\mu\mathcal{R}_M}{\mu\mathcal{R}_M1},$$

where for bounded measurable f ,

$$\mathcal{R}_M f(x) = \int_0^\infty dt \mathbb{E}_x[f(Y_t)1\{\tau_M > t\}].$$

We have that

$$\begin{aligned} \pi_M - \pi &= \frac{(\mu\mathcal{R}1)\mu\mathcal{R}_M - (\mu\mathcal{R}_M1)\mu\mathcal{R}}{(\mu\mathcal{R}_M1)(\mu\mathcal{R}1)} \\ &= \frac{\mu\mathcal{R}1(\mu\mathcal{R}_M - \mu\mathcal{R}) + (\mu\mathcal{R}1 - \mu\mathcal{R}_M1)\mu\mathcal{R}}{(\mu\mathcal{R}_M1)(\mu\mathcal{R}1)}. \end{aligned}$$

So we would like to bound

$$|\mu\mathcal{R}f - \mu\mathcal{R}_M f|$$

for arbitrary bounded measurable f .

For any nonnegative bounded measurable f ,

$$\begin{aligned}
 |\mu\mathcal{R}f - \mu\mathcal{R}_Mf| &\leq \int \mu(dx) \int dt |\mathbb{E}_x[f(Y_t)1\{\tau_M > t\}1\{\tau_M^e > t\}] \\
 &\quad - \mathbb{E}_x[f(X_t)1\{\tau_M > t\}]| \\
 &\leq \int \mu(dx) \int dt \mathbb{E}_x[f(Y_t)(1 - 1\{\tau_M^e > t\})1\{\tau_M > t\}] \\
 &\leq \|f\|_\infty \int \mu(dx) \int dt \mathbb{E}_x[1\{\tau_M^e \leq t, \tau_M > t\}] \\
 &= \|f\|_\infty \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t, \tau_M > t).
 \end{aligned}$$

Since we are assuming that we have a lower bound $\underline{\kappa}$ on the regeneration rate, and Assumption 27 holds, we can stochastically bound $\tau_M \leq \tau'$ where $\tau' \sim \text{Exp}(\underline{\kappa})$ and τ' is independent of everything else. So continuing the chain of inequalities,

$$\begin{aligned}
 &\leq \|f\|_\infty \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t, \tau' > t) \\
 &= \|f\|_\infty \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) \mathbb{P}_x(\tau' > t) \\
 &= \|f\|_\infty \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}}.
 \end{aligned}$$

A universal upper bound on this quantity is $\|f\|_\infty/\underline{\kappa}$.

For for a given continuous nonnegative bounded f with $\|f\|_\infty \leq 1$ we get the following bounds:

$$\begin{aligned}
 |\pi_M f - \pi f| &\leq \frac{\mu\mathcal{R}1\|f\|_\infty \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}} + \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}} |\mu\mathcal{R}f|}{(\mu\mathcal{R}_M1)(\mu\mathcal{R}1)} \\
 &\leq \frac{\int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}}}{\mu\mathcal{R}_M1} + \frac{\int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}}}{\mu\mathcal{R}_M1} \\
 &= \frac{2 \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}}}{\mu\mathcal{R}_M1} \\
 &\leq \frac{2 \int \mu(dx) \int dt \mathbb{P}_x(\tau_M^e \leq t) e^{-t\underline{\kappa}}}{\mu\mathcal{R}1}.
 \end{aligned}$$

Since this bound is valid for only nonnegative bounded f , in order to bound $\|\pi_M - \pi\|_1$ we pick up an additional factor of 2.

This concludes the proof of Theorem 30.

Acknowledgments. We would like to thank Martin Kolb for supplying arguments pertaining to domains of self-adjoint operators. We would also like to thank Jere Koskela, Divakar Kumar, Hector McKimm and Abraham Ng for interesting discussions on aspects of this paper. We would like to thank the Associate Editor and anonymous referee for their comments which have substantially improved the paper.

The first author was supported by EPSRC CoSiNES Grant EP/R034710/1 and by EPSRC OxWaSP CDT through Grant EP/L016710/1. Most of his work was conducted as a doctoral student at the Department of Statistics, University of Oxford, and a preliminary version of this work appears in his doctoral thesis, Wang (2020).

The second author was supported by The Alan Turing Institute under the EPSRC Grant EP/N510129/1, and his contribution was largely conducted while at the University of Warwick.

The third author was supported by EPSRC Grants EP/R034710/1, EP/R018561/1, EP/K014463/1.

The fourth author was supported by BBSRC Grant BB/S001824/1.

REFERENCES

- ADAMS, R. A. (1975). *Sobolev Spaces. Pure and Applied Mathematics* **65**. Academic Press, New York. [MR0450957](#)
- ASMUSSEN, S. (2003). *Applied Probability and Queues: Stochastic Modelling and Applied Probability*, 2nd ed. *Applications of Mathematics (New York)* **51**. Springer, New York. [MR1978607](#)
- ASMUSSEN, S. and GLYNN, P. W. (2007). *Stochastic Simulation: Algorithms and Analysis. Stochastic Modelling and Applied Probability* **57**. Springer, New York. [MR2331321](#)
- BARBOUR, A. D. and POLLETT, P. K. (2010). Total variation approximation for quasi-stationary distributions. *J. Appl. Probab.* **47** 934–946. [MR2752899](#) <https://doi.org/10.1017/s0021900200007270>
- BARBOUR, A. D. and POLLETT, P. K. (2012). Total variation approximation for quasi-equilibrium distributions, II. *Stochastic Process. Appl.* **122** 3740–3756. [MR2965923](#) <https://doi.org/10.1016/j.spa.2012.07.004>
- BARTLETT, M. S. (1960). *Stochastic Population Models in Ecology and Epidemiology. Methuen's Monographs on Applied Probability and Statistics*. Methuen & Co., Ltd., London. [MR0118550](#)
- BENAIM, M., CLOEZ, B. and PANLOUP, F. (2018). Stochastic approximation of quasi-stationary distributions on compact spaces and applications. *Ann. Appl. Probab.* **28** 2370–2416. [MR3843832](#) <https://doi.org/10.1214/17-AAP1360>
- BIERKENS, J., FEARNHEAD, P. and ROBERTS, G. (2019). The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Statist.* **47** 1288–1320. [MR3911113](#) <https://doi.org/10.1214/18-AOS1715>
- BOUCHARD-CÔTÉ, A., VOLLMER, S. J. and DOUCET, A. (2018). The bouncy particle sampler: A non-reversible rejection-free Markov chain Monte Carlo method. *J. Amer. Statist. Assoc.* **113** 855–867. [MR3832232](#) <https://doi.org/10.1080/01621459.2017.1294075>
- BROCKWELL, A. E. and KADANE, J. B. (2005). Identification of regeneration times in MCMC simulation, with application to adaptive schemes. *J. Comput. Graph. Statist.* **14** 436–458. [MR2161623](#) <https://doi.org/10.1198/106186005X47453>
- CAPUTO, P. and QUATTROPANI, M. (2019). Mixing time of PageRank surfers on sparse random digraphs. Preprint. Available at [arXiv:1905.04993](https://arxiv.org/abs/1905.04993).
- COLLET, P., MARTÍNEZ, S. and SAN MARTÍN, J. (2013). *Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems. Probability and Its Applications (New York)*. Springer, Heidelberg. [MR2986807](#) <https://doi.org/10.1007/978-3-642-33131-2>
- DARROCH, J. N. and SENETA, E. (1965). On quasi-stationary distributions in absorbing discrete-time finite Markov chains. *J. Appl. Probab.* **2** 88–100. [MR0179842](#) <https://doi.org/10.2307/3211876>
- DAVIS, M. H. A. (1984). Piecewise-deterministic Markov processes: A general class of nondiffusion stochastic models. *J. Roy. Statist. Soc. Ser. B* **46** 353–388. [MR0790622](#)
- DEMUTH, M. and VAN CASTEREN, J. A. (2000). *Stochastic Spectral Theory for Selfadjoint Feller Operators: A Functional Integration Approach. Probability and Its Applications*. Birkhäuser, Basel. [MR1772266](#) <https://doi.org/10.1007/978-3-0348-8460-0>
- DEVROYE, L. (1986). *Nonuniform Random Variate Generation*. Springer, New York. [MR0836973](#) <https://doi.org/10.1007/978-1-4613-8643-8>
- DOOB, J. L. (1945). Markoff chains—Denumerable case. *Trans. Amer. Math. Soc.* **58** 455–473. [MR0013857](#) <https://doi.org/10.2307/1990339>
- DURMUS, A., GUILLIN, A. and MONMARCHÉ, P. (2018). Piecewise deterministic Markov processes and their invariant measure. Preprint. Available at [arXiv:1807.05421](https://arxiv.org/abs/1807.05421).
- ETHIER, S. N. and KURTZ, T. G. (1986). *Markov Processes: Characterization and Convergence. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics*. Wiley, New York. [MR0838085](#) <https://doi.org/10.1002/9780470316658>
- GILKS, W. R., ROBERTS, G. O. and SAHU, S. K. (1998). Adaptive Markov chain Monte Carlo through regeneration. *J. Amer. Statist. Assoc.* **93** 1045–1054. [MR1649199](#) <https://doi.org/10.2307/2669848>
- GRIMMETT, G. R. and STIRZAKER, D. R. (2001). *Probability and Random Processes*, 3rd ed. Oxford Univ. Press, New York. [MR2059709](#)

- HOBERT, J. P., JONES, G. L., PRESNELL, B. and ROSENTHAL, J. S. (2002). On the applicability of regenerative simulation in Markov chain Monte Carlo. *Biometrika* **89** 731–743. MR1946508 <https://doi.org/10.1093/biomet/89.4.731>
- JACOB, P. E., O'LEARY, J. and ATCHADÉ, Y. F. (2020). Unbiased Markov chain Monte Carlo methods with couplings. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **82** 543–600. MR4112777 <https://doi.org/10.1111/rssb.12336>
- KUMAR, D. (2019). On a stochastic approximation and quasi-stationary approach for the simulation of Bayesian posterior distributions, with application to tall data. Ph.D. thesis, Univ. Warwick.
- LEE, A., DOUCET, A. and LATUSZYŃSKI, K. (2014). Perfect simulation using atomic regeneration with application to sequential Monte Carlo. Preprint. Available at [arXiv:1407.5770](https://arxiv.org/abs/1407.5770).
- MEYN, S. P. and TWEEDIE, R. L. (1993). *Markov Chains and Stochastic Stability. Communications and Control Engineering Series*. Springer, London. MR1287609 <https://doi.org/10.1007/978-1-4471-3267-7>
- MINH, D. L., MINH, D. D. L. and NGUYEN, A. L. (2012). Regenerative Markov chain Monte Carlo for any distribution. *Comm. Statist. Simulation Comput.* **41** 1745–1760. MR2924015 <https://doi.org/10.1080/03610918.2011.615433>
- MOYAL, J. E. (1957). Discontinuous Markoff processes. *Acta Math.* **98** 221–264. MR0093824 <https://doi.org/10.1007/BF02404475>
- MURDOCH, D. J. (2000). Exact sampling for Bayesian inference: Unbounded state spaces. In *Monte Carlo Methods (Toronto, ON, 1998)*. *Fields Inst. Commun.* **26** 111–121. Amer. Math. Soc., Providence, RI. MR1772310
- MURDOCH, D. J. and GREEN, P. J. (1998). Exact sampling from a continuous state space. *Scand. J. Stat.* **25** 483–502. MR1650023 <https://doi.org/10.1111/1467-9469.00116>
- MYKLAND, P., TIERNEY, L. and YU, B. (1995). Regeneration in Markov chain samplers. *J. Amer. Statist. Assoc.* **90** 233–241. MR1325131
- NUMMELIN, E. (1978). A splitting technique for Harris recurrent Markov chains. *Z. Wahrsch. Verw. Gebiete* **43** 309–318. MR0501353 <https://doi.org/10.1007/BF00534764>
- POLLOCK, M., FEARNHEAD, P., JOHANSEN, A. M. and ROBERTS, G. O. (2020). Quasi-stationary Monte Carlo methods and the ScaLE algorithm. *J. Roy. Statist. Soc. Ser. B*. To appear.
- PROPP, J. G. and WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9** 223–252. MR1611693 [https://doi.org/10.1002/\(SICI\)1098-2418\(199608/09\)9:1/2<223::AID-RSA14>3.3.CO;2-R](https://doi.org/10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.3.CO;2-R)
- SHARPE, M. (1988). *General Theory of Markov Processes. Pure and Applied Mathematics* **133**. Academic Press, Boston, MA. MR0958914
- THORISSON, H. (2000). *Coupling, Stationarity, and Regeneration. Probability and Its Applications (New York)*. Springer, New York. MR1741181 <https://doi.org/10.1007/978-1-4612-1236-2>
- TIERNEY, L. (1996). Introduction to general state-space Markov chain theory. In *Markov Chain Monte Carlo in Practice. Interdiscip. Statist.* 59–74. CRC Press, London. MR1397968
- VANETTI, P., BOUCHARD-CÔTÉ, A., DELIGIANNIDIS, G. and DOUCET, A. (2017). Piecewise deterministic Markov chain Monte Carlo. Preprint. Available at [arXiv:1707.05296](https://arxiv.org/abs/1707.05296).
- WANG, A. Q. (2020). Theory of killing and regeneration in continuous-time Monte Carlo sampling. Ph.D. thesis, Univ. Oxford.
- WANG, A. Q., ROBERTS, G. O. and STEINSALTZ, D. (2020). An approximation scheme for quasi-stationary distributions of killed diffusions. *Stochastic Process. Appl.* **130** 3193–3219. MR4080743 <https://doi.org/10.1016/j.spa.2019.09.010>
- WANG, A. Q. and STEINSALTZ, D. (2019). A note on the jump locations of Markov processes. Preprint. Available at [arXiv:1901.07321](https://arxiv.org/abs/1901.07321).
- WANG, A. Q., KOLB, M., ROBERTS, G. O. and STEINSALTZ, D. (2019). Theoretical properties of quasi-stationary Monte Carlo methods. *Ann. Appl. Probab.* **29** 434–457. MR3910008 <https://doi.org/10.1214/18-AAP1422>