A Discrete Bouncy Particle Sampler *

C. Sherlock[1] and A.H. Thiery[2]

[1]Department of Statistics, Lancaster University, United Kingdom
[2]Department of Statistics & Applied Probability, National University of Singapore

Abstract

Most Markov chain Monte Carlo methods operate in discrete time and are reversible with respect to the target probability. Nevertheless, it is now understood that the use of non-reversible Markov chains can be beneficial in many contexts. In particular, the recently-proposed Bouncy Particle Sampler leverages a continuous-time and non-reversible Markov process and empirically shows state-of-the-art performances when used to explore certain probability densities; however, its implementation typically requires the computation of local upper bounds on the gradient of the log target density.

We present the Discrete Bouncy Particle Sampler, a general algorithm based upon a guided random walk, a partial refreshment of direction, and a delayed-rejection step. We show that the Bouncy Particle Sampler can be understood as a scaling limit of a special case of our algorithm. In contrast to the Bouncy Particle Sampler, implementing the Discrete Bouncy Particle Sampler only requires point-wise evaluation of the target density and its gradient. We propose extensions of the basic algorithm for situations when the exact gradient of the target density is not available. In a Gaussian setting, we establish a scaling limit for the radial process as dimension increases to infinity. We leverage this result to obtain the theoretical efficiency of the Discrete Bouncy Particle Sampler as a function of the partial-refreshment parameter, which leads to a simple and robust tuning criterion. A further analysis in a more general setting suggests that this tuning criterion applies more generally. Theoretical and empirical efficiency curves are then compared for different targets and algorithm variations.

Key words: Markov Chain Monte-Carlo; Non-reversible Samplers; Bouncy Particle Sampler; Scaling Limit.

1 Introduction

Markov Chain Monte Carlo (MCMC) algorithms provide Monte Carlo approximations to expectations with respect to a given probability distribution, π, via an ergodic Markov chain whose invariant distribution is π. Non-reversible Markov Chaine Monte Carlo samplers, of which the Hamiltonian Monte Carlo algorithm (Duane et al., 1987) is perhaps one of the most successful and

*This is the author accepted manuscript; the version of record is available from https://academic.oup.com/biomet/advance-article/doi/10.1093/biomet/asab013/6151695

widely-used examples, are known to enjoy desirable mixing properties in several contexts. Indeed, several theoretical results quantify the advantages of non-reversible samplers. For example, Diaconis et al. (2000) obtains rates of convergence for a non-reversible version of the random walk algorithm; subsequently and inspired by Diaconis et al. (2000), Chen et al. (1999) describes the best acceleration achievable through the idea of lifting. On a different note, Hwang et al. (2015), Lelièvre et al. (2013), Rey-Bellet and Spiliopoulos (2015) and Duncan et al. (2016) investigate and quantify the advantages offered by leveraging (a discretization of) a non-reversible diffusion process for computing Monte-Carlo averages: in many settings, it can be proved that the standard reversible Langevin dynamics is the worst in terms of asymptotic variances among a large class of diffusion processes that are ergodic with respect to a given target distribution. More recently, different designs of non-reversible MCMC sampler have been proposed. The Zig-Zag sampler, an instance of the large class of Piecewise-Deterministic-Markov-Processes, first obtained as a scaling limit of a lifted Metropolis–Hastings Markov chain (Bierkens and Roberts, 2017), is a continuous-time non-reversible Markov process that can be used for computing ergodic averages, and can be used for efficiently exploring Bayesian posterior distributions in the Big-Data regime (Bierkens et al., 2019). Inspired from the physics literature (Peters and de With, 2012), the Bouncy Particle Sampler (Bouchard-Côté et al., 2017) is another continuous-time non-reversible Monte-Carlo sampler that demonstrates state-of-the-art performance when used to explore certain Bayesian posterior distributions. Fearnhead et al. (2018) reviews the Zig-Zag sampler and the Bouncy Particle Sampler and describes some extensions. The Bouncy Particle Sampler or the Zig-Zag Sampler requires more than simple point-wise evaluations of the log-target density and its gradient: one typically needs local upper bounds on derivatives of the log-target density. Unfortunately, those bounds are unavailable or difficult to compute in many applied situations. Consequently, such continuous-time samplers cannot be directly used in these settings. This article presents a discrete-time MCMC sampler, inspired by the Bouncy Particle Sampler, that can be implemented when only point-wise evaluations of the target-density and its gradient are available.

Our algorithm, the Discrete Bouncy Particle Sampler, is described in detail in Section 2.1. It extends the statespace from a position to a position and a direction in the same way that the Zig-Zag and Bouncy Particle samplers do; however, our algorithm operates in discrete time and is based upon the guided random walk of Gustafson (1998). The guided random walk combines two reversible kernels to create a non-reversible kernel which heads in a specific direction until a rejection occurs, at which point it reverses direction. Our key addition is a particular delayed-rejection proposal (Tierney and Mira, 1999), used after any initial rejection, potentially avoiding many inefficient direction reversals. The delayed-rejection move is analogous to the bounce in the Bouncy Particle Sampler and we show that the Discrete Bouncy Particle Sampler can be viewed as a time discretization of the Bouncy Particle Sampler. An alternative discretization of the Bouncy Particle Sampler, based on the reflective slice sampler (Neal, 2003) is described and extended in the independent work of Vanetti et al. (2017). Importantly, several interesting extensions have recently been proposed (Vanetti et al., 2017; Wu and Robert, 2017; Wu and Robert, 2020; Monmarché, 2019; Michel et al., 2020) to scale and enhance this class of Piecewise-Deterministic-Markov-Processes MCMC samplers.

As with the Bouncy Particle Sampler, our algorithm can be reducible. To solve this issue, we perturb the direction vector at the end of every iteration. The size of this per-iteration perturbation has a substantial impact on the performance of the algorithm, in a similar way to the occasional, complete direction refresh of the Bouncy Particle Sampler (Bierkens et al., 2018). Analogously

to the independent investigations for the Bouncy Particle Sampler in Bierkens et al. (2018), for the Discrete Bouncy Particle Sampler exploring a Gaussian target we use diffusion-approximation arguments to describe the limit of the radial process as dimension increases to infinity. We then leverage this to obtain the theoretical efficiency of the Discrete Bouncy Particle Sampler as a function of the partial-refreshment parameter, κ. This leads to a simple and robust tuning criterion that is described in Section 3.2. In more practical developments, we show that a surrogate may be substituted for the gradient of the target density when it is computationally expensive or impossible to obtain. Our final contribution is a construction that allows the user to choose to only calculate a fixed number of orthogonal components of the gradient vector.

Throughout the remainder of this article, the distribution of interest is referred to as $\pi(\mathsf{d}x)$, and is assumed to have support on the standard Euclidean space $\mathcal{X} \equiv (\mathbb{R}^d, \langle \cdot, \cdot \rangle)$ with associated norm $\|\cdot\|$. This distribution is assumed to have a density with respect to Lebesgue measure, which, with a slight abuse of notation, is also denoted by π. We set $x \wedge y = \min(x, y)$. For a distribution π and a π-integrable function φ, the quantity $\pi(\varphi)$ refers to the expectation of φ under π. For $x \in \mathbb{R}$, its positive and negative part are denoted as $x_+ \geq 0$ and $x_- \geq 0$ so that $x = x_+ - x_-$. For two vectors $u, v \in \mathbb{R}^d$, their dot product is $\langle u, v \rangle = u_1 \, v_1 + \ldots + u_d \, v_d$. For any time t, t^- refers to the instant in time just prior to t and t^+ to that just after t.

2 The Discrete Bouncy Particle Sampler

2.1 Algorithm description

The Discrete Bouncy Particle Sampler operates on the extended state space $\mathcal{X} \times \mathcal{S}$, where $\mathcal{S} \subseteq \mathcal{X}$, and explores the extended target distribution

$$\widetilde{\pi}(\mathsf{d}x, \mathsf{d}u) = \pi(\mathsf{d}x) \otimes \rho(\mathsf{d}u)$$

where $\rho(\mathsf{d}u)$ is an auxiliary spherically symmetric distribution with support $\mathcal{S} \subset \mathcal{X}$. Section 2.2 describes several standard choices of auxiliary distributions. Henceforth, we will refer to the variable $x \in \mathcal{X}$ as the position of a particle and the variable $u \in \mathcal{S}$ as its direction. The bounce after which the Discrete Bouncy Particle Sampler is named enters through the operator $u \mapsto \mathscr{R}_v(u)$ that reflects the vector $u \in \mathcal{S}$ with respect to the hyperplane orthogonal to the vector $v \in \mathcal{X} \setminus \{0\}$,

$$\mathscr{R}_v(u) \equiv u - 2 \, \frac{\langle u, v \rangle}{\|v\|^2} \, v. \tag{1}$$

For any vector $v \in \mathcal{X} \setminus \{0\}$, the reflection operator $\mathscr{R}_v$ is an involution $\mathscr{R}_v \circ \mathscr{R}_v(u) = u$ that preserves norms. This remark underlies the proof of correctness of the Discrete Bouncy particle Sampler whose details are presented in the Supplementary Material. As discussed in Michel et al. (2020), it is possible to rely on more general reflection operators. Most of the methods developed in this text extends to these variants, although we concentrate on Equation (1) for ease of exposition. The Discrete Bouncy Particle Sampler relies on a non-vanishing vector field

$$\mathscr{F} : \mathcal{X} \to \mathcal{X} \setminus \{0\}.$$

In practice, this vector field is either chosen as $\mathscr{F}(x) = \nabla \log \pi(x)$, replaced by an arbitrary modification when the gradient vanishes, or as an approximation of it, as described in Section 2.4.

The quantity $\mathscr{R}_{\mathscr{F}(x)}(u)$ represents the resulting direction when a particle with incoming direction u performs an elastic bounce off the hyperplane orthogonal to the vector $\mathscr{F}(x)$.

The Discrete Bouncy Particle Sampler deterministically cycles through two Markovian transitions that leave the extended target distribution $\widetilde{\pi}$ invariant, (1) a Position Update, with a possible Direction Reflection (2) a Direction Refreshment. The resulting scheme is, in general, non-reversible. For the direction refreshment, one can choose any Markov transition kernel $\mathsf{P}_\rho(u, \mathsf{d}u')$ that leaves the auxiliary distribution ρ invariant. For a discretization parameter $\delta > 0$, and a current state $(x_k, u_k) \in \mathcal{X} \times \mathcal{S}$, the algorithm proceeds as follows.

1. POSITION UPDATE: Generate a proposal $(x', u') = (x_k + \delta\, u_k, u_k)$. With position update probability

$$\alpha_{\mathrm{pu}}(x_k, u_k) \equiv 1 \wedge \frac{\pi(x')}{\pi(x_k)}, \tag{2}$$

   set $(\hat{x}_k, \hat{u}_k) = (x', u')$ and go to Step 3. Otherwise, proceed to Step 2.

2. DIRECTION REFLECTION: consider $u'' = \mathscr{R}_{\mathscr{F}(x')}(u')$ and $x'' = x' + \delta\, u''$. With direction reflection probability

$$\alpha_{\mathrm{dr}}(x_k, u_k) \equiv 1 \wedge \left\{ \frac{1 - \alpha_{\mathrm{pu}}(x'', -u'')}{1 - \alpha_{\mathrm{pu}}(x_k, u_k)} \times \frac{\pi(x'')}{\pi(x_k)} \right\}, \tag{3}$$

   set $(\hat{x}_k, \hat{u}_k) = (x'', u'')$. Otherwise, negate the direction by setting $(\hat{x}_k, \hat{u}_k) = (x_k, -u_k)$.

3. DIRECTION REFRESHMENT: Set $(x_{k+1}, u_{k+1}) = (\hat{x}_k, U)$ where $\mathrm{P}(U \in A) = \mathsf{P}_\rho(\hat{u}_k, A)$.

By construction, the Direction Refreshment step preserves the extended target distribution. The fact that the combination of the Position Update and Direction Reflection steps also preserves the extended distribution is discussed in the Supplementary Material. This can be understood as a slight generalization of the standard delayed rejection mechanism (Tierney and Mira, 1999) when applied to a deterministic and volume preserving proposal. A similar scheme was proposed independently in Vanetti et al. (2017). Furthermore, and importantly for Section 2.4, the algorithm remains valid if the deterministic vector field is replaced by a randomized version of it. The proof of correctness is identical to the deterministic case and is briefly discussed in the Supplementary Material.

A simple thought experiment, such as considering a target density π with spherically-symmetric contours and $\mathsf{P}_\rho(u, \mathsf{d}u') = \delta_u(\mathsf{d}u')$, shows that, as with the Bouncy Particle Sampler, the Discrete Bouncy Particle Sampler can be reducible. The choice and tuning of the direction refreshment operator P_ρ is consequently important in practice and is discussed at length in the sequel.

2.2 Direction dynamics

The choice of the auxiliary distribution $\rho(\mathsf{d}u)$ and the update operator P_ρ, which we now allow to depend on a tuning refreshment parameter, κ, has a major influence on the efficiency of the resulting Discrete Bouncy Particle Sampler. Two standard choices for the isotropic distribution ρ are the uniform distribution ρ_S on the unit sphere of $\mathbb{R}^d$ and the centred Gaussian distribution ρ_G with covariance $(1/d)\, I_d$. The scaling of the covariance matrix implies that $\int \|u\|^2\, \rho_G(\mathsf{d}u) =$

$\int \|u\|^2 \rho_S(\mathsf{d}u) = 1$, which ensures that the effect on the algorithm of the tuning parameters δ and κ are comparable across the different auxiliary distributions. As will be described in Section 2.3, it is convenient to think of $\mathsf{P}_\rho(u, \mathsf{d}u')$ as the discretization between time t and $t+\delta$ of a continuous-time Markov process $\{V_t\}_{t\geq 0}$ that leaves $\rho(\mathsf{d}u)$ invariant, $\mathsf{P}_\rho(u, \mathsf{d}u') = \mathrm{P}(V_{t+\delta} \in \mathsf{d}u' | V_t = u)$. We now describe several standard ways of generating a random variable U that, conditioned on a current direction $u \in \mathcal{S}$, is approximately distributed as $\mathsf{P}_\rho(u, \mathsf{d}u')$.

- FULL REFRESH: for $\rho = \rho_S$ or $\rho = \rho_G$ and an update rate $\kappa > 0$ the Markov process with generator $\mathcal{L}^{(V)}\varphi(u) = \kappa\,(\rho(\varphi) - \varphi(u))$ completely refreshes the direction at rate κ and has a mixing time of $\mathcal{O}(1/\kappa)$. For a time discretization parameter $0 < \delta < 1$, set

$$U = B_\kappa^\delta\, u + (1 - B_\kappa^\delta)\, \xi$$

  where $\xi \sim \rho$ and B_κ^δ is a Bernoulli random variable with $\mathrm{P}(B_\kappa^\delta = 1) = \exp(-\kappa\,\delta)$.

- ORNSTEIN-UHLENBECK REFRESH: for $\rho = \rho_G$ and an update rate $\kappa > 0$, the Ornstein-Uhlenbeck process $\mathsf{d}V_t = -(\kappa/2)\, V_t\, \mathsf{d}t + (\kappa/d)^{1/2}\, \mathsf{d}W$ leaves ρ_G invariant and has a mixing time of $\mathcal{O}(1/\kappa)$. Set

$$U = \alpha\, u + (1 - \alpha^2)^{1/2}\, \xi \tag{4}$$

  for $\xi \sim \rho_G$ and $\alpha = \exp(-\kappa\,\delta/2)$.

- BROWNIAN MOTION ON THE UNIT SPHERE: for $\rho = \rho_S$ and an update rate $\kappa > 0$, consider the Brownian motion on the unit sphere in $\mathbb{R}^d$. It is described by the stochastic differential equation

$$\mathsf{d}V_t = -\frac{\kappa}{2}\, V_t\, \mathsf{d}t + \{\kappa/(d-1)\}^{1/2}\, P_\perp(V_t)\, \mathsf{d}W \tag{5}$$

  where W is a standard Brownian motion in $\mathbb{R}^d$ and $P_\perp(V_t) \in \mathbb{R}^{d,d}$ is the orthogonal projection on the hyperplane orthogonal to V_t. The Brownian motion on the unit sphere leaves ρ_S invariant and has a mixing time of $\mathcal{O}(1/\kappa)$. One can define U by normalizing an Ornstein-Uhlenbeck update (4),

$$U = \frac{\alpha\, u + (1 - \alpha^2)^{1/2}\, \xi}{\|\alpha\, u + (1 - \alpha^2)^{1/2}\, \xi\|},$$

  for $\xi \sim \rho_G$ and $\alpha = \exp(-\kappa\,\delta/2)$.

2.3 Continuous-time limit

In this section we show that, as $\delta \to 0$, the Discrete Bouncy Particle Sampler converges to a well defined continuous-time and piecewise-continuous Markov process. This result clarifies the role of the reflection operator $\mathscr{R}_{\mathscr{F}(x)}$ and explains the connection between the Discrete Bouncy Particle Sampler and the continuous time Bouncy Particle Sampler. For any time discretization parameter $\delta > 0$, consider a Discrete Bouncy Particle Sampler chain $\{(x_k^\delta, u_k^\delta)\}_{k\geq 0}$ with update operator $\mathsf{P}_\rho^\delta(u, \mathsf{d}u')$. The superscript δ indicates, as will be made clearer in Assumption $A3$ stated

below, that the update operator $\mathsf{P}_\rho^\delta(u, \mathsf{d}u')$ is obtained as the discretization of a Markov process $\{V_t\}_{t\geq 0}$ between time t and $t+\delta$. For a time horizon $T > 0$, consider the continuous time process $\overline{z}_t^\delta = (\overline{x}_t^\delta, \overline{u}_t^\delta)$ defined on the interval $[0, T]$ by setting

$$\overline{z}_{k\delta}^\delta = (\overline{x}_{k\delta}^\delta, \overline{u}_{k\delta}^\delta) = (x_k^\delta, u_k^\delta) \in \mathcal{X} \times \mathcal{S}$$

for any integer $0 \leq k \leq T/\delta$ and linearly interpolation in between. The main result of this section, Proposition 1, whose proof is presented in the Supplementary Material, relies on the following regularity assumptions.

A1 The function $x \mapsto \log \pi(x)$ is twice differentiable with a bounded second derivative.

A2 The vector field $\mathscr{F} : \mathcal{X} \to \mathcal{X} \setminus \{0\}$ is continuous.

A3 There exists a continuous time Markov process $\{V_t\}_{t\geq 0}$ with generator $\mathcal{L}^{(V)}$ such that, for any time discretization parameter $\delta > 0$, the transition kernel P_ρ^δ describes the transition of the Markov process V in the sense that $\mathsf{P}_\rho^\delta(u, \mathsf{d}u') = \mathrm{P}(V_{t+\delta} \in \mathsf{d}u' | V_t = u)$. We assume that the trajectories of the Markov process $\{V_t\}_{t\geq 0}$ are almost surely continuous.

Proposition 1. *Let Assumptions A(1-2-3) hold and consider a fixed time horizon $T > 0$. As $\delta \to 0$, the sequence of continuous time processes $\overline{z}_t^\delta = (\overline{x}_t^\delta, \overline{u}_t^\delta)$ converges weakly in the Skorokhod topology to the bivariate Markov process $\overline{Z}_t = (\overline{X}_t, \overline{U}_t)$ with generator*

$$\begin{aligned} \mathcal{L}^{(\overline{Z})}\,\varphi(x,u) &= \mathcal{L}^{(V)}\varphi(x,u) + \langle u, \nabla_x\,\varphi(x,u)\rangle \\ &\quad + \lambda(x,u)\left([\mathcal{A}(x,u)\,\varphi\{x, \mathscr{R}_{\mathscr{F}(x)}(u)\} + \{1-\mathcal{A}(x,u)\}\,\varphi(x,-u)] - \varphi(x,u)\right) \end{aligned} \tag{6}$$

with rate $\lambda(x,u) \equiv \langle -\nabla \log \pi(x), u\rangle_+$ and acceptance probability

$$\mathcal{A}(x,u) \equiv 1 \wedge \frac{\lambda\{x, -\mathscr{R}_{\mathscr{F}(x)}(u)\}}{\lambda(x,u)} \in [0,1]. \tag{7}$$

The limiting Markov process $\overline{Z}_t = (\overline{X}_t, \overline{U}_t)$ with generator (6) evolves according to the dynamics

$$\mathsf{d}\overline{X}_t = \overline{U}_t\,\mathsf{d}t, \quad \mathsf{d}\overline{U}_t = \mathsf{d}V_t, \tag{8}$$

in between events that arrive at rate $\lambda(\overline{X}_t, \overline{U}_t)$. When such an event is triggered, the direction is reflected, i.e. $\overline{U}_t = \mathscr{R}_{\mathscr{F}(X_{t^-})}(\overline{U}_{t^-})$, with probability $\mathcal{A}(\overline{X}_{t^-}, \overline{U}_{t^-})$, and completely reversed, i.e. $\overline{U}_t = -\overline{U}_{t^-}$, with probability $1 - \mathcal{A}(\overline{X}_{t^-}, \overline{U}_{t^-})$. Possible choices of Markovian dynamics with generator $\mathcal{L}^{(V)}$ in Assumption A2 are detailed in Section 2.2. In the case when $\mathscr{F}(x) = \nabla \log \pi(x)$ and $\mathcal{L}^{(V)}\varphi(u) = \kappa\,\{\rho(\varphi) - \varphi(u)\}$, for a fixed refreshment rate $\kappa > 0$, the limiting Markov process is the standard Bouncy Particle Sampler (Bouchard-Côté et al., 2017). The interested reader is referred to Vanetti et al. (2017); Wu and Robert (2017) for other interesting generalizations

In order to understand the influence of the vector field $\mathscr{F}$, it is instructive to study the limiting acceptance probability (7). The limiting process $\overline{Z}$ is rejection free, i.e. never backtracks, if for any $(x,u) \in \mathcal{X} \times \mathcal{S}$ we have that $\lambda\{x, -\mathscr{R}_{\mathscr{F}(x)}(u)\} = \lambda(x,u)$. It is readily seen that this condition is equivalent to choosing $\mathscr{F}(x)$ proportional to $\nabla \log \pi(x)$ for any $x \in \mathcal{X}$ where this quantity does not vanish. In other words, any other choice of vector field $\mathscr{F}$ leads to a limiting process that is not rejection-free. Section 2.4 describes ways to efficiently approximate this optimal choice when evaluating $\nabla \log \pi(x)$ is not computationally efficient.

2.4 Approximate reflections

As described in the previous section, vector fields that lead to a rejection-free algorithm in the limit $\delta \to 0$ are such that $\mathscr{F}(x)$ is proportional to $\nabla \log \pi(x)$ for all $x \in \mathcal{X}$ where this quantity is non-zero. When computing the gradient of the log-density is not computationally feasible, one can instead use a vector field $\mathscr{F}$ that only approximates $\nabla \log \pi$, necessarily paying the price of a non-zero probability for the limiting algorithm to backtrack. Another strategy, similar to the one presented in Fielding et al. (2011), consists in choosing $\mathscr{F}(x)$ as the gradient of an approximate surrogate target distribution.

Assume that the Discrete Bouncy Particle Sampler stands at $(x, u) \in \mathcal{X} \times \mathcal{S}$ and that a Direction Reflection is attempted. The exact gradient $g(x) \equiv \nabla \log \pi(x)$ being unavailable, one can instead numerically evaluate $g(x) \in \mathbb{R}^d$ along a set of $n_{\text{cpt}} \leq d$ random directions described by a set of n_{cpt} mutually orthogonal unit vectors $\underline{\zeta} \equiv (\zeta_1, \dots, \zeta_{n_{\text{cpt}}})$ generated from a distribution $G_x(\mathrm{d}\underline{\zeta})$ that may depend on the current position $x \in \mathcal{X}$ but not the current direction $u \in \mathcal{S}$. As described in the Supplementary Material, the fact that the distribution $G_x(\mathrm{d}\underline{\zeta})$ does not depend on the current direction ensures that the resulting algorithm has the correct invariant distribution. A standard choice consists in orthonormalizing a set of n_{cpt} vectors generated from an isotropic Gaussian distribution. The approximate gradient is then defined as $\widetilde{g}(x) = \sum_i^{n_{\text{cpt}}} \langle g(x), \zeta_i \rangle \, \zeta_i$, where each coefficient$\langle g(x), \zeta_i \rangle$ can be evaluated numerically (Ramm and Smirnova, 2001). In other words, $\widetilde{g}(x)$ is the orthogonal projection of the vector $g(x)$ onto the plane $V(\underline{\zeta}) \equiv \mathrm{span}(\zeta_1, \dots, \zeta_{n_{\text{cpt}}})$. Decomposing the direction as $u = u^\perp + u^\parallel$ with $u^\parallel \in V(\underline{\zeta})$ and $u^\perp \in V(\underline{\zeta})^\perp$, the following two modified bounce operators $u \mapsto u'$ both lead to a correct algorithm:

1. The updated direction $u' \in \mathcal{S}$ can be defined as the reflection with respect to the hyperplane orthogonal to $\widetilde{g}(x)$, i.e. $u' = \mathscr{R}_{\widetilde{g}(x)}(u)$. This update can also be expressed as
$$u' = \mathscr{R}_{\widetilde{g}(x)}(u) = u^\perp + \mathscr{R}_{\widetilde{g}(x)}(u^\parallel).$$

2. One can also completely reflect the component of $u \in \mathcal{S}$ that is orthogonal to the plane $V(\underline{\zeta})$. In other words, the updated direction u' is defined as
$$u' = -u^\perp + \mathscr{R}_{\widetilde{g}(x)}(u^\parallel). \tag{9}$$

While both reflection operators are valid, we have empirically found that the operator (9) leads to better mixing properties.

2.5 Preconditioning

A general target may have very different length scales in different directions and, just as with Metropolis-Hastings algorithms such as the random walk Metropolis or the Metropolis-Adjusted Langevin algorithm (Roberts and Rosenthal, 2001), the efficiency can be improved, often by several orders of magnitude, by preconditioning. Consider an invertible matrix $\Gamma \in \mathbb{R}^{d,d}$ and define the whitened variable $\widetilde{x}$ implicitly defined as $x = \Gamma \widetilde{x}$. Instead of using the Discrete Bouncy Particle Sampler for exploring a target density $\pi(x)$, one can instead explore the whitened density $\tilde{\pi}(\widetilde{x}) \propto \pi(\Gamma \widetilde{x})$. See Pakman et al. (2017) for an analogous description of preconditioning for the BPS. Typically, the transformation $\widetilde{x} \mapsto \Gamma \widetilde{x}$ is chosen such that the transformed density $\tilde{\pi}$ is as isotropic as possible. A standard strategy is to choose Γ such that $\Gamma\Gamma^\top$ is an approximation of the covariance matrix of the target distribution π, or of the negative inverse Hessian of $\log \pi$ at a mode.

3 Algorithm tuning through diffusion approximation

3.1 Diffusion Limit

In this subsection, we derive a diffusion approximation, in a Gaussian setting, for the log-target process in the high-dimensional regime $d \to \infty$, as the discretization parameter $\delta \to 0$. We are interested in understanding the mixing properties of the process $k \mapsto \log \pi^d(x_k^{d,\delta})$ when $\{(x_k^{d,\delta}, u_k^{d,\delta})\}_{k\geq 0}$ is a Discrete Bouncy Particle Sampler chain targeting the d-dimensional Gaussian distribution π^d. One could also study the mixing properties of any other functional of the Markov chain: for a fixed dimension $D \geq 1$ and a sequence of functions $F_d : \mathbb{R}^d \to \mathbb{R}^D$ indexed by $d \geq 1$, one can investigate the properties of the stochastic process $k \mapsto F_d(x_k^{d,\delta})$. For example, Deligiannidis et al. (2021) proves scaling limits of a finite and fixed set of coordinates, which corresponds to the projection operator $F_d(x_1, \ldots, x_d) = (x_1, \ldots, x_D)$. Our choice $F^d(x) = \log \pi^d(x)$ is motivated by the empirical observation that in many scenarios, when the Discrete Bouncy Particle Sampler is employed, the mixing of the log-target process is much slower than the mixing of any given coordinate. Similar empirical observations are discussed in Terenin and Thorngren (2018), and Bierkens et al. (2018) proves different diffusion limits for the standard Bouncy Particle Sampler when full refreshments are used to ensure ergodicity. Using entirely different techniques, Andrieu et al. (2021) obtains insights into the scaling of the Bouncy Particle Sampler and more general Piecewise Deterministic Markov Processes. The expression we obtain for the limiting process allows us in Section 3.2 to formulate robust strategies for tuning the refreshment parameter $\kappa > 0$. Consider a Discrete Bouncy Particle Sampler with reflection $\mathscr{R}_{\mathscr{F}(x)}$ where $\mathscr{F}(x) = \nabla \log \pi^d$ and when the refreshment updates are obtained as discretization of a Brownian motion on the unit sphere, as described in Section 2.2. We concentrate on the case where the target distribution π_d is a centred d-dimensional Gaussian density with covariance $\sigma^2 \, \mathrm{I}_d$, for a fixed standard deviation $\sigma > 0$. Thanks to the rotational symmetry of π_d, all the bounce attempts are accepted and studying the log-target process is equivalent to studying the radial process $k \mapsto \|x_k^{d,\delta}\|$. Proposition 1 identifies, for a fixed dimension $d \geq 1$, the scaling limit as $\delta \to 0$ of the Discrete Bouncy Particle Sampler chain $\{x_k^{d,\delta}, u_k^{d,\delta}\}_{k\geq 0}$. For concreteness, we denote this limiting continuous-time Markov process, whose generator is described in Proposition 1, as $(\overline{X}_t^d, \overline{U}_t^d)$. Now, for investigating the high-dimensional regime $d \to \infty$, note that for a sequence of random variables $X^d \sim \pi_d$, the sequence $\|X^d\| - \sigma \, d^{1/2}$ converges in distribution towards a centred Gaussian distribution with variance $\sigma^2/2$. We consequently define the shifted process

$$R_t^d = \|\overline{X}_{d\times t}^d\| - \sigma \, d^{1/2}. \tag{10}$$

Note that, in the definition of the process R_t^d, time has been accelerated by a factor of d. Proposition 2 stated below shows that, in order to observe a non-degenerate scaling limit as $d \to \infty$, this acceleration factor is the correct one. As will be demonstrated, the mixing properties of R_t^d are closely related to the mixing properties of the scalar jump-diffusion $\{\theta_t^\kappa\}_{t\geq 0}$ with generator

$$\mathcal{L}^{(\kappa,\sigma)} = \frac{\kappa}{2}\mathcal{L}^K + \frac{1}{\sigma}\mathcal{L}^J \qquad \text{where} \qquad \begin{cases} \mathcal{L}^{(K)}\varphi(\theta) = -\theta\,\varphi'(\theta) + \varphi''(\theta) \\ \mathcal{L}^{(J)}\varphi(\theta) = \varphi'(\theta) + \theta_+ \{\varphi(-\theta) - \varphi(\theta)\}. \end{cases} \tag{11}$$

The operator $\mathcal{L}^{(K)}$ is the generator of an Ornstein-Uhlenbeck process that is reversible with the standard Gaussian density. Similarly, $\mathcal{L}^{(J)}$ is the generator of the Markov process with unit drift

and reflections $\theta \mapsto -\theta$ that occur at rate $\theta_+ = \max(0, \theta)$. It can readily be checked that this process also leaves the standard Gaussian distribution invariant. Combining these two facts show that the process θ_t^κ also leaves the standard Gaussian distribution invariant. We denote by $V_\sigma(\kappa)$ the asymptotic variance of ergodic averages along θ_t^κ defined as

$$V_\sigma(\kappa) = \lim_{T \to \infty} \text{Var}\left(\frac{1}{\sqrt{2\,T}} \int_0^T \theta_t^\kappa \, \mathsf{d}t\right). \tag{12}$$

There is no closed form expression for the quantity $V_\sigma(\kappa)$ but it can easily be approximated numerically, as displayed in Figure 1. Note that $V_\sigma(\kappa) \to 0$ as $\kappa \to 0$ and $\kappa \to \infty$. Proposition 2, whose proof can be found in the Supplementary Material, shows that the asymptotic variance $V_\sigma(\kappa)$ dictates the mixing rate of the log-target process. The higher the asymptotic variance $V_\sigma(\kappa)$, the faster the mixing of the radial process.

Proposition 2. *Let $T > 0$ be a finite time horizon. For a fixed refreshment parameter $\kappa > 0$, the sequence of accelerated processes R^d defined in Equation* (10) *converges weakly as $d \to \infty$ in $C([0, T], \mathbb{R})$ to the Ornstein-Uhlenbeck process*

$$d\overline{R}_t = -\frac{V_\sigma(\kappa)}{\sigma^2/2}\, \overline{R}_t \, dt + \{2\, V_\sigma(\kappa)\}^{1/2}\, dW. \tag{13}$$

The velocity function $V_\sigma(\kappa)$ is defined in Equation (12)*.*

The process (13) is an Ornstein-Uhlenbeck that is reversible with respect to the centred Gaussian distribution with variance $\sigma^2/2$. Since the Ornstein-Uhlenbeck (13) has a mixing time of order $\mathcal{O}(1)$, this indicates that in the high-dimensional regime $d \to \infty$ and $\delta \to 0$, one can expect (Roberts and Rosenthal, 2016) the log-target process to mix on a time scale of order $\mathcal{O}(d/\delta)$. When implementing the Discrete Bouncy Particle Sampler in practice, the parameter δ should be chosen small enough to guarantee that the acceptance rate remains high-enough, but not smaller. Similar guidelines for the Hamiltonian Monte-Carlo method are described in Betancourt et al. (2014). The optimal tuning of the parameter δ, and study of its dependence with respect to the dimensionality of the target distribution, is beyond the scope of this article. Instead, we concentrate on the tuning of the refreshment parameter κ. When optimising the mixing of the log-target process, we observe empirically that the tuning of the parameter κ is insensitive to the value of δ. This is in part because whatever the value of δ, as long as it is sufficiently small, the log-target process is an approximation of the limiting diffusion (13) (see also Section 3.2); empirical evidence that this insensitivity continues to hold for large δ is provided in Section 4.1.

3.2 Tuning of the refreshment parameter κ

The limiting diffusion (13) obtained in Section 3.1 indicates that, for an isotropic Gaussian target distribution with marginal variance σ^2 and in the regime $d \to \infty$, optimising the efficiency of the Discrete Bouncy Particle Sampler is achieved by choosing a refreshment parameter $\kappa > 0$ that maximizes the velocity $V_\sigma(\kappa)$. In other words, for a given marginal variance $\sigma^2 > 0$, the optimal refreshment rate $\kappa_\star(\sigma)$ is given by

$$\kappa_\star(\sigma) = \mathsf{argmax}\, \{\kappa \mapsto V_\sigma(\kappa)\}.$$

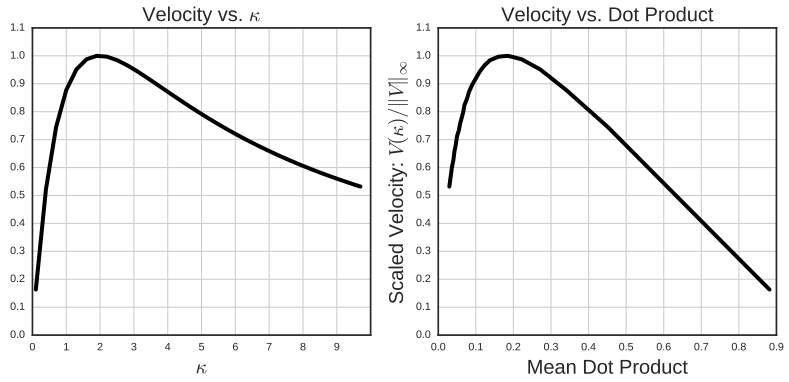

Figure 1: Velocity $V(\kappa)$ for $\sigma = 1$ as a function of $\kappa > 0$ (Left) and as a function of the mean dot product (Right). The scale of the velocity function being irrelevant, only the rescaled velocity function $\kappa \mapsto V(\kappa)/\|V\|_\infty$ is depicted.

Furthermore, a change of time argument immediately shows that $\kappa_\star(\sigma) = \kappa_\star/\sigma$ with $\kappa_\star \equiv \kappa_\star(\sigma = 1)$. In practice, the variance parameter $\sigma^2 > 0$ is not known so that the optimal refreshment parameter is not directly accessible. To make progress, denote by $\{\tau_j\}_{j\geq 1}$ the (strictly increasing) sequence of time indices at which Direction Reflection events are attempted (and always accepted in the Gaussian setting). We denote by $u_{\tau_j^-} \in \mathcal{S}$ the direction right before a reflection event, and by $u_{\tau_j^+} \in \mathcal{S}$ the direction right after the reflection. For tuning purposes, we propose to monitor the dot product $\beta \in [-1, 1]$ between the direction vectors right after and before the Direction Reflection attempts,

$$\beta = \left\langle u_{\tau_j^+}, u_{\tau_{j+1}^-} \right\rangle. \tag{14}$$

For a Discrete Bouncy Particle Sampler evolving at stationarity, as $\delta \to 0$ and for any fixed dimension $d \geq 2$, consider the distribution $\mu^d(\mathsf{d}\beta; \kappa, \sigma)$ of these dot products. One can readily check that if $\{x_k^{d,\delta}, u_k^{d,\delta}\}_{k\geq 0}$ is Discrete Bouncy Particle Sampler chain with parameters $\kappa, \delta > 0$ exploring the centred d-dimensional Gaussian with marginal standard deviation σ then, for any scaling factor $s > 0$, the Markov chain defined as $\{s \times x_k^{d,\delta}, u_k^{d,\delta}\}_{k\geq 0}$ is also Discrete Bouncy Particle Sampler chain, with refreshment parameter κ/s and time discretization parameter $s\,\delta > 0$, exploring the centred d-dimensional Gaussian with marginal standard deviation $s\,\sigma$. It follows that $\mu^d(\mathsf{d}\beta, \kappa, \sigma) = \mu^d(\mathsf{d}\beta; \kappa/s, s\,\sigma)$ for any scaling factor $s > 0$. Consequently, since $\kappa_\star(\sigma) = \kappa_\star/\sigma$, the distribution $\mu^d(\mathsf{d}\beta; \kappa_\star(\sigma), \sigma) \equiv \mu_\star^d(\mathsf{d}\beta)$ does not depend on the standard deviation σ. It is straightforward to numerically estimate the average dot product at optimality,

$$\beta_\star \equiv \lim_{d\to\infty} \int_{-1}^{1} \beta\, \mu_\star^d(\mathsf{d}\beta) \approx 0.2.$$

Figure 1 illustrates this optimality result. Very low values $\beta \ll \beta_\star$ indicate that the directions are updated too frequently, leading to an inefficient random-walk behaviour. High values $\beta \approx 1$ indicate that the directions are not updated frequently enough, leading to an inefficient exploration of the state space. The case $\beta = 1$ corresponds to the case when the direction are not updated at

all, which is known in the Gaussian case to lead to a reducible Markov Chain with an incorrect invariant distribution. For tuning the refreshment parameter $\kappa > 0$ of a general Discrete Bouncy Particle Sampler, we consequently propose to estimate empirically the expectation at stationarity of the quantity β in (14). Let $\{\tau_j\}_{j\geq 1}$ now denote the realization of the sequence of indices at which direction reflection attempts occur. A direction reflection attempt is accepted with probability (3), otherwise the direction is negated. We define

$$\widehat{\beta} = \frac{1}{J}\sum_{j=1}^{J}\left\langle u_{\tau_j^+}, u_{\tau_{j+1}^-}\right\rangle$$

and choose $\kappa > 0$ so that this quantity approximately equals its optimal value $\beta_\star \approx 0.2$. Just as with the estimation of acceptance rates when tuning the scaling of various algorithms (Roberts and Rosenthal, 2001), and unlike the Effective Sample Size itself that is notoriously difficult to reliably estimate, the mean dot product can be estimated accurately from short MCMC runs. Importantly, and as described in Section 4, we have found this tuning procedure to be robust with respect to departure from Gaussianity and to approximately hold in non-isotropic and relatively low-dimensional settings with $\delta \gg 0$.

3.3 Non-isotropic target and non-zero δ

The diffusion limit in Section 3.1 was obtained as $\delta \downarrow 0$ and for an isotropic Gaussian target where the direction reflection proposals are always accepted. In this Section, we consider the non-isotropic case of a d-dimensional target distribution defined as

$$\pi^{(d)}(x^{(d)}) = \prod_{i=1}^{d} \gamma_i \exp\{f(\gamma_i x_i)\}, \tag{15}$$

for inverse scalings $\gamma_i > 0$ drawn independently from a distribution with second moment $\mathrm{m}_2(\gamma) \equiv \mathrm{E}(\gamma^2)$ and finite third moment. For a Discrete Bouncy Particle Sampler Markov chain exploring this density, denote by $\alpha_{\mathrm{pu}}^{(d)}\{x^{(d)}, u^{(d)}\}$ and $\alpha_{\mathrm{dr}}^{(d)}\{x^{(d)}, u^{(d)}\}$ the relevant acceptance probabilities in dimension $d \geq 1$. The corresponding acceptance rates are $\alpha_{\mathrm{pu}}^{(d)} \equiv \mathrm{E}[\alpha_{\mathrm{pu}}^{(d)}\{X^{(d)}, U^{(d)}\}]$ and $\alpha_{\mathrm{dr}}^{(d)} \equiv \mathrm{E}[\alpha_{\mathrm{dr}}^{(d)}\{X^{(d)}, U^{(d)}\}|\mathsf{DR}]$, where $\{X^{(d)}, U^{(d)}\}$ follows its stationary distribution, and DR is the event that the initial proposal is rejected and that a Direction Reflection is attempted. The definition of the acceptance rate $\alpha_{\mathrm{dr}}^{(d)}$ requires conditioning on there being a direction update in order to ensure that the quantity $\alpha_{\mathrm{dr}}^{(d)}\{X^{(d)}, U^{(d)}\}$ is well-defined. Theorem 1, whose proof is in the Supplementary Material, shows that, in the high-dimensional regime $d \to \infty$ and for a fixed $\delta > 0$, the probability of accepting the direction reflection proposals converges towards one.

Theorem 1. *Let $\{X_t^{(d)}\}_{t=0}^{\infty}$ be a d-dimensional Discrete Bouncy Particle Sampler Markov chain created by the algorithm described in Section 2.1, exploring the target density defined in Equation* (15)*. Assume that the function $f : \mathbb{R} \to \mathbb{R}$ has a Lipschitz second-derivative and that the quantity $J \equiv \mathrm{E}\{f'(\xi)^2\} = -\mathrm{E}\{f''(\xi)\}$, for a random variable ξ with density $\exp\{f(\xi)\}$, is finite. The acceptance rates are such that*

$$\lim_{d\to\infty} \alpha_{\mathrm{pu}}^{(d)} = 2\,\Phi\left[-\frac{\delta}{2}\left\{\mathrm{m}_2(\gamma)\,J\right\}^{1/2}\right] \quad \textit{and} \quad \lim_{d\to\infty} \alpha_{\mathrm{dr}}^{(d)} = 1$$

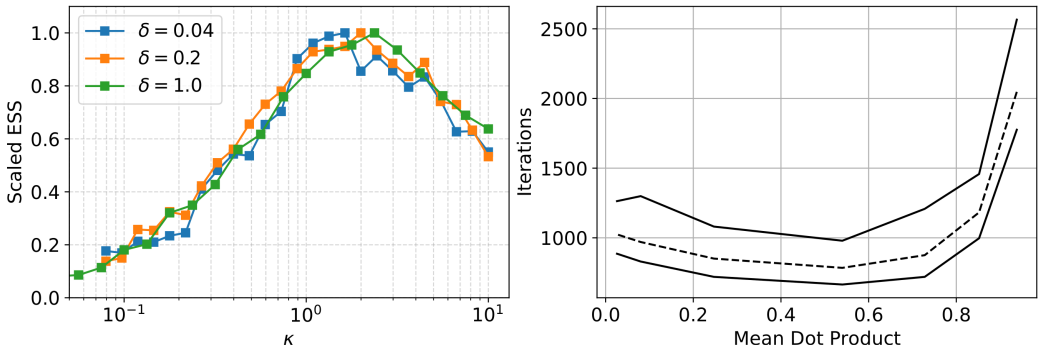

Figure 2: Left: relative effective sample size against κ for a 100-dimensional standard Gaussian target. Right: maximum (solid), minimum (solid) and median (dashed) iterations to converge from a random tail point of target (16) as a function of the mean dot product statistic β in a stationary run.

where $\Phi(t) = (2\pi)^{-1/2} \int_0^t e^{-t^2/2}\, dt$ *is the standard Gaussian cumulative function.*

In the Supplementary Material, a simulation study verifies the theorem for a particular target distribution and a variety of time discretization parameters δ. This study also suggests that, if δ is kept fixed with respect to the dimension, then $1 - \alpha_{\mathrm{dr}}^{(d)} \sim 1/d$. In Section 2.2, the refreshment parameter κ was intentionally defined so that as $\delta \downarrow 0$ the effect of κ, which is on the mixing time of the velocity refreshment, should be insensitive to the time discretization paremeter δ. Given the definitions of κ and δ, it is not unreasonable to expect this insensitivity to carry through to macroscopic values of δ. Section 4.1 details a simulation study that confirms that the tuning choice for κ is insensitive to the value of δ. This, together with Theorem 1, suggests that the tuning strategy obtained from the isotropic Gaussian diffusion limit could be applicable for more general high-dimensional targets.

4 Simulation studies

4.1 Tuning of κ is insensitive to δ

We first[1] investigate the robustness to the choice of δ of the effect of κ on efficiency. We considered an isotropic Gaussian target distribution with dimension $d = 100$ and, for a range of $\delta \in \{0.04, 0.2, 1.0\}$, we estimated the efficiency of the Discrete Bouncy Particle Sampler as a function of the refreshment parameter κ. For a fixed δ, the relative variation in $1 - \alpha_{\mathrm{pu}}$ was small, varying less than 5% from a central value over the whole range of κ values. Central values of $1 - \alpha_{\mathrm{pu}}$ were 2% ($\delta = 0.04$), 8% ($\delta = 0.2$) and 38% ($\delta = 1.0$). These values span the range of potential interest in the applications we have looked at. The left panel of Figure 2 suggests that the optimal value of the refreshment parameter κ is insensitive to the time discretization parameter δ. This insensitivity over more than an order of magnitude of δ values suggests that the value of the

[1]Code is available at this GitHub repository

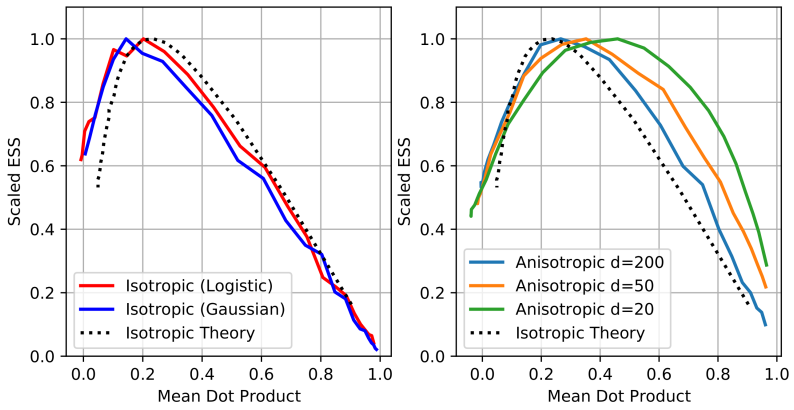

Figure 3: Scaled effective sample size as a function of the mean dot product for several values of κ, obtained from MCMC runs of length $N = 10^6$. The Discrete Bouncy Particle Sampler is applied to the target distributions π^{log} and π^{iso} (Left) and π^{aniso} using $d = 20$, $d = 50$ and $d = 200$ (Right)

optimal refreshment parameter κ obtained in the limit $\delta \to 0$ and, hence, $1 - \alpha_{\text{pu}} \to 0$, is indicative of the optimal refreshment parameter κ when δ is macroscopic and $1 - \alpha_{\text{pu}} \gg 0$.

4.2 Robustness of advice to departures from Gaussianity and isotropy

We next investigate the robustness of our tuning advice to departures of the target from Gaussianity and isotropy.

Consider three scenarios: an isotropic multivariate logistic distribution π^{logis} with density $\prod_{i=1}^{d} \exp(x_i)/[1 + \exp(x_i)]^2$, an isotropic Gaussian distribution π^{iso} with density proportional to $\prod_{i=1}^{d} \exp(-x_i^2/2)$ and a non-isotropic multivariate Gaussian distribution π^{aniso} with density proportional to $\prod_{i=1}^{d} \exp\{-x_i^2/(2\,\sigma_i^2)\}$. In this section, we choose $d = 100$ for both isotropic targets, and $d \in \{20, 50, 200\}$ for the anisotropic target. In order to test the robustness of our tuning guidelines to non-isotropic distributions, we chose the scales $\sigma_1 < \ldots < \sigma_d$ linearly separated between $\sigma_1 = 1$ and $\sigma_d = 10$. The scaled effective sample size curves as a function of the mean dot-product are displayed in Figure 3. There is extremely good agreement with the theory developed in Section 4 for the isotropic distribution π^{iso} and the approximately isotropic distribution π^{logistic}. Not surprisingly, mild departure from the theory is observed for strongly non-isotropic distributions such as π^{aniso}. However, for dimension $d = 200$ the departure, especially in terms of the optimal dot product, is barely noticeable. When $d = 20$ and $d = 50$, however, our proposed guideline, *i.e.* tune the refreshment rate κ such that the mean dot-product $\beta \approx 0.2$, leads only to a loss of efficiency of approximately 10% and 5% respectively.

4.3 Convergence and tail behaviour

One of the most commonly used algorithms for inference in high-dimensional scenarios is Hamiltonian Monte Carlo (Duane et al., 1987). However, it is well known (Livingstone et al., 2019) that due to the dependence of the Leapfrog step on $\|\nabla \log \pi\|$, Hamiltonian Monte Carlo is not geometrically ergodic on targets with tails lighter than those of a Gaussian. By contrast, the Discrete

Bouncy Particle Sampler depends on $\nabla \log \pi$ only through the equivalent normalized vector. The Supplementary Material details a simulation study on a non-isotropic, light-tailed target where a tuned Hamiltonian Monte Carlo algorithm is nearly 50% more efficient than a tuned discrete bouncy particle sampler when started from stationarity. However, with the same tunings, but when started from a random point in the tail of the distribution, Hamiltonian Monte Carlo does not even move, whereas the discrete bouncy particle sampler quickly converges to the centre of the posterior mass.

Whilst our theory combined with empirical verifications suggests that the optimal refreshment rate κ at stationarity is that which leads to an average dot product $\beta \approx 0.2$, it may be that a different value is optimal for convergence from the tails. Finally, therefore, we examine the effect of κ, when the Discrete Bouncy Particle Sampler is used to explore the target distribution in $\mathbb{R}^d$ with $d = 50$ and a density of

$$f(x) \propto \exp\left(-\frac{1}{a}\|x\|_{\mathrm{M}}^a\right) \qquad \text{for} \quad \|x\|_{\mathrm{M}}^2 \equiv \sum_{i=1}^{d} \frac{x_i^2}{\sigma_i^2} \tag{16}$$

where $a = 4$ and $\sigma_1, \ldots, \sigma_d$ are as described in Section 4.2. In $d > 1$ dimensions, the modal value for $\|x\|_{\mathrm{M}}$ is $r_\star \equiv (d-1)^{1/a}$. For several values of $\kappa > 0$, we repeated the following experiment $n = 100$ times: sample a random unit vector $z \in \mathbb{S}^{d-1}$, set the initial position $x_0 = 10\, r_\star \cdot (\sigma_1 z_1, \ldots, \sigma_d z_d) \in \mathbb{R}^d$ far out in the tail of the distribution, run the Discrete Bouncy Particle Sampler with $\delta = 2.0$ (a sensible value to explore the body of the target) until $\|x\|_{\mathrm{M}} \leq r_\star$ and note the iteration number at which that happened. The right-hand panel of Figure 2 shows the minimum, maximum and median iteration number at which convergence (by this measure) was achieved. Except for very low κ values, which lead to large dot product statistics, the behaviour is robust to the choice of κ, suggesting that it is reasonable to apply our tuning criterion when the chain is started away from its stationary distribution.

4.4 The Markov modulated Poisson process

A simulation study on the eight-dimensional posterior of a non-trivial statistical model, the Markov modulated Poisson process, is detailed in the Supplementary Material. We find that mixing efficiency of $\log \pi$ is optimized at a dot product statistic of $\beta \approx 0.4$, tuning to $\beta \approx 0.2$ would lead to only a 10% reduction in efficiency. When either preconditioning or using only $n_{\text{cpt}} = 3$ random components of the gradient vector, the optimal efficiency is achieved for a dot product statistics of $\beta \approx 0.2$.

5 Discussion

The key advantage of the Discrete Bouncy Particle Sampler over its continuous-time counterpart is that posterior and gradient evaluations can be treated as a "black box" with no requirement to bound the gradient so as to apply Poisson thinning. Unlike the continuous-time algorithm, there is a chance that an attempted bounce will be rejected and the particle will approximately backtrack, however for sensible tunings, the back-tracking probability converges towards zero as the dimension of the problem increases.

The average computational cost per iteration is insensitive to the choice of κ since the direction is updated every iteration, and κ has little effect on the number of steps between potential bounces. Thus it is sufficient for the purposes of this article to describe and empirically record efficiency in terms of effective sample size rather than effective samples per unit of time. The only exception to this is when we compare against Hamiltonian Monte Carlo in Appendix C.1.

The dot-product tuning diagnostic maps κ to an absolute scale via the properties of the posterior, just as the acceptance rate diagnostic does for the scaling in the random-walk Metropolis algorithm. When $\kappa = 0$ the velocity direction just before the next bounce is identical to that just after the previous bounce, and the dot product is unity. The mixing time of the refreshment process is $1/\kappa$; when this is small compared with the time between bounces or, equivalently, the length scale of the target, the two velocity directions bear little relation to each other, and the dot product is small.

Whilst this article offers theory-based practical advice on the tuning of the refreshment parameter, κ, it does not tackle the choice of the discretization parameter, δ. In contrast to the insensitivity of computational cost to the choice of κ, increasing δ increases the frequency of potential bounces and hence of expensive gradient calculations, and this would need to be accounted for in any analysis.

Acknowledgements

Work by CS was supported by EPSRC grant EP/P033075/1. AHT acknowledges support from the Singapore Ministry of Education Tier 2 Grant (MOE2016-T2-2-135) and a Young Investigator Award Grant (NUSYIA FY16 P16; R-155-000-180-133).

References

Andrieu, C., Durmus, A., Nüsken, N., and Roussel, J. (2021). Hypocoercivity of piecewise deterministic Markov process-Monte Carlo. *Annals of Applied Probability*. To Appear.

Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M., Stuart, A., et al. (2013). Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534.

Betancourt, M., Byrne, S., and Girolami, M. (2014). Optimizing the integrator step size for Hamiltonian Monte Carlo. *arXiv preprint arXiv:1411.6669*.

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(3):1288–1320.

Bierkens, J., Kamatani, K., and Roberts, G. O. (2018). High-dimensional scaling limits of piecewise deterministic sampling algorithms. *arXiv preprint arXiv:1807.11358*.

Bierkens, J. and Roberts, G. (2017). A piecewise deterministic scaling limit of lifted Metropolis–Hastings in the Curie–Weiss model. *The Annals of Applied Probability*, 27(2):846–882.

Bouchard-Côté, A., Vollmer, S. J., and Doucet, A. (2017). The bouncy particle sampler: A nonreversible rejection-free Markov chain Monte Carlo method. *Journal of the American Statistical Association*.

Chen, F., Lovász, L., and Pak, I. (1999). Lifting Markov chains to speed up mixing. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 275–281. ACM.

Deligiannidis, G., Paulin, D., and Doucet, A. (2021). Randomized Hamiltonian Monte Carlo as scaling limit of the bouncy particle sampler and dimension-free convergence rates. *Annals of Applied Probability*. accepted.

Diaconis, P., Holmes, S., and Neal, R. M. (2000). Analysis of a nonreversible Markov chain sampler. *Annals of Applied Probability*, pages 726–752.

Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics letters B*, 195(2):216–222.

Duncan, A. B., Lelievre, T., and Pavliotis, G. (2016). Variance reduction using nonreversible Langevin samplers. *Journal of statistical physics*, 163(3):457–491.

Fearnhead, P., Bierkens, J., Pollock, M., and Roberts, G. O. (2018). Piecewise-deterministic Markov processes for continuous-time Monte Carlo. *Statist. Sci.*, 33(3):386–412.

Fearnhead, P. and Sherlock, C. (2006). An exact Gibbs sampler for the Markov-modulated Poisson process. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(5):767–784.

Fielding, M., Nott, D. J., and Liong, S.-Y. (2011). Efficient MCMC schemes for computationally expensive posterior distributions. *Technometrics*, 53(1):16–28.

Gustafson, P. (1998). A guided walk Metropolis algorithm. *Statistics and Computing*, 8(4):357–364.

Hwang, C.-R., Normand, R., and Wu, S.-J. (2015). Variance reduction for diffusions. *Stochastic Processes and their Applications*, 125(9):3522–3540.

Lelièvre, T., Nier, F., and Pavliotis, G. A. (2013). Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion. *Journal of Statistical Physics*, 152(2):237–274.

Livingstone, S., Betancourt, M., Byrne, S., and Girolami, M. (2019). On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli*, 25(4A):3109–3138.

Michel, M., Durmus, A., and Sénécal, S. (2020). Forward event-chain Monte Carlo: Fast sampling by randomness control in irreversible Markov chains. *Journal of Computational and Graphical Statistics*, pages 1–14.

Monmarché, P. (2019). Kinetic walks for sampling. *arXiv preprint arXiv:1903.00550*.

Neal, R. M. (2003). Slice sampling. *Ann. Statist.*, 31(3):705–767.

Pakman, A., Gilboa, D., Carlson, D., and Paninski, L. (2017). Stochastic bouncy particle sampler. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2741–2750, International Convention Centre, Sydney, Australia. PMLR.

Papanicolaou, G. (1977). Martingale approach to some limit theorems. In *Papers from the Duke Turbulence Conference, Duke Univ., Durham, NC, 1977*.

Pavliotis, G. and Stuart, A. (2008). *Multiscale methods: averaging and homogenization*. Springer Science & Business Media.

Peters, E. A. and de With, G. (2012). Rejection-free Monte Carlo sampling for general potentials. *Physical Review E*, 85(2):026703.

Ramm, A. and Smirnova, A. (2001). On stable numerical differentiation. *Mathematics of computation*, 70(235):1131–1153.

Rey-Bellet, L. and Spiliopoulos, K. (2015). Irreversible Langevin samplers and variance reduction: a large deviations approach. *Nonlinearity*, 28(7):2081.

Roberts, G. O., Gelman, A., Gilks, W. R., et al. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, 7(1):110–120.

Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, 16(4):351–367.

Roberts, G. O. and Rosenthal, J. S. (2016). Complexity bounds for Markov chain Monte Carlo algorithms via diffusion limits. *Journal of Applied Probability*, 53(2):410–420.

Terenin, A. and Thorngren, D. (2018). A piecewise deterministic Markov process via radius-angle swaps in hyperspherical coordinates. *arXiv preprint arXiv:1807.00420*.

Tierney, L. and Mira, A. (1999). Some adaptive Monte Carlo methods for Bayesian inference. *Statistics in Medicine*, 18:2507–2515.

Vanetti, P., Bouchard-Côté, A., Deligiannidis, G., and Doucet, A. (2017). Piecewise-deterministic Markov chain Monte Carlo. *arXiv preprint arXiv:1707.05296*.

Weinan, E. (2011). *Principles of multiscale modeling*. Cambridge University Press.

Wu, C. and Robert, C. P. (2017). Generalized Bouncy Particle Sampler. *ArXiv e-prints*.

Wu, C. and Robert, C. P. (2020). The coordinate sampler: A non-reversible Gibbs-like MCMC sampler. *Statistics and Computing*, 30:721–730.

A Correctness and proofs of propositions

A.1 Correctness of the Discrete Bouncy Particle Sampler

In this section, we prove the correctness of a slightly more general version of the Discrete Bouncy Particle Sampler than the one described in the main text. This added generality is exploited in Section 2.4. Consider a generalized reflection operator $\mathsf{R} : \mathbb{R}^d \times \mathbb{R}^d$ such that for every $x \in \mathcal{X} \equiv \mathbb{R}^d$, the mapping $u \mapsto \mathsf{R}(u, x)$ satisfies the following three conditions:

B1 For any $u \in \mathbb{R}^d$, we have that $\mathsf{R}(-\mathsf{R}(u,x),x) = -u$

B2 The mapping $u \mapsto \mathsf{R}(u,x)$ is volume preserving.

B3 The mapping $u \mapsto \mathsf{R}(u,x)$ preserves norms, $\|\mathsf{R}(u,x)\| = \|u\|$.

For $(x_k, u_k) \in \mathbb{R}^d \times \mathbb{R}^d$ and a time discretization parameter $\delta > 0$, consider the Markov kernel $(x_k, u_k) \mapsto (x_{k+1}, u_{k+1})$ defined as the composition of the following three steps.

Step 1. Generate a proposal $(x', u') = (x_k + \delta\, u_k, -u_k)$. With probability

$$\alpha_{\mathrm{pu}}(x_k, u_k) \equiv 1 \wedge \frac{\pi(x')}{\pi(x_k)},$$

set $(\hat{x}_k, \hat{u}_k) = (x', -u')$ and go to Step 3. Otherwise, proceed to Step 2.

Step 2. consider $u'' = -\mathsf{R}(u, x')$ and $x'' = x' - \delta\, u''$. With probability

$$\alpha_{\mathsf{R}}(x_k, u_k) \equiv 1 \wedge \left\{ \frac{1 - \alpha_{\mathrm{pu}}(x'', u'')}{1 - \alpha_{\mathrm{pu}}(x_k, u_k)} \times \frac{\pi(x'')}{\pi(x_k)} \right\},$$

set $(\hat{x}_k, \hat{u}_k) = (x'', u'')$. Otherwise, set $(\hat{x}_k, \hat{u}_k) = (x_k, u_k)$.

Step 3. Reverse the direction: $(x_{k+1}, u_{k+1}) = (\hat{x}_k, -\hat{u}_k)$

Lemma 1. *Consider any spherically symmetric probability density $\rho(u)$. Under Assumptions B1-2-3, the Markov kernel described by Step 1-2-3 leaves the density $\widetilde{\pi}(x,u) = \pi(x)\,\rho(u)$ invariant.*

Proof. Since ρ is spherically symmetric and R preserves norms, Step 3 leaves $\widetilde{\pi}$ invariant. Now, the combination of Step 1-2 is exactly a Delayed Rejection (Tierney and Mira, 1999) Markov kernel with two proposal mechanisms: $(x,u) \mapsto (x + \delta\, u, -u)$ and $(x,u) \mapsto (x + \delta\, u + \delta\, \mathsf{R}(u, x + \delta u), -\mathsf{R}(u, x + \delta u))$ and target density $\pi(x)\,\rho(x)$. Algebra directly shows that these two proposals are volume preserving involutions. To conclude the proof of the lemma, it thus suffices to show that the usual delayed rejection scheme for sampling from a density $\mu(z)$ on a state-space $\mathcal{Z} \subset \mathbb{R}^N$ remains valid with deterministic proposals $z \mapsto z' = T_1(z)$ and $z \mapsto z'' = T_2(z)$ that are volume preserving involutions. For an arbitrary bounded test function $\varphi : \mathcal{Z} \to \mathbb{R}$, one needs to show that

$$\int \left[\varphi\left\{T_1(z)\right\}\, \alpha_1(z) + \varphi\left\{T_2(z)\right\}\, \left\{1 - \alpha_1(z)\right\}\, \alpha_2(z) + \varphi(z)\, \alpha_3(z)\right]\, \mu(z)\, \mathsf{d}z = \int \varphi(z)\, \mu(z)\, \mathsf{d}z$$

with $\alpha_1(z) = 1 \wedge \mu\{T_1(z)\}/\mu(z)$ and $\alpha_3(z) = 1 - \alpha_1(z) - \{1 - \alpha_1(z)\}\,\alpha_2(z)$ and

$$\alpha_2(z) = 1 \wedge \left\{ \frac{1 - \alpha_1\{T_2(z)\}}{1 - \alpha_1(z)} \times \frac{\mu\{T_2(z)\}}{\mu(z)} \right\}.$$

Algebra shows that this is equivalent to proving that

$$\begin{aligned} &\int [\varphi\{T_1(z)\} - \varphi(z)] \times [\mu(z) \wedge \mu\{T_1(z)\}] \\ &\quad + \int [\varphi\{T_2(z)\} - \varphi(z)] \times [\mu(z)\{1 - \alpha_1(z)\} \wedge \mu\{T_2(z)\}(1 - \alpha_1\{T_2(z)\})]\, \mathsf{d}z = 0. \end{aligned} \tag{17}$$

Since T_1 and T_2 are involutions that preserve volume, a change of variable $z \mapsto T_1(z)$ shows that the first integral in Equation (17) also equals its negation, and hence vanishes. And similarly, the change of variable $z \mapsto T_2(z)$ shows that the second integral in Equation (17) also vanishes. This concludes the proof of the lemma. □

In Section 2.1, the combination of the Position Update and Direction Update is equivalent to Step 1-2-3 with the operator $\mathsf{R}(u, x) = \mathscr{R}_{\mathscr{F}(x)}(u)$. Since algebra shows that the Conditions B1-2-3 are satisfied, Lemma 1 thus shows the correctness of the Discrete Bouncy Particle Sampler as described in Section 2.1.

Indeed, one can consider mixtures of operators that satisfy Conditions B1-2-3. Namely, for a conditional probability distribution $\mathrm{M}(x, \mathrm{d}\omega)$ and an operator $\widetilde{\mathsf{R}} : \mathbb{R}^d \times \mathbb{R}^d \times \Omega \to \mathbb{R}^d$ such that for any value of $(x, \omega) \in \mathbb{R}^d \times \Omega$ the operator $u \mapsto \widetilde{\mathsf{R}}(u, x, \omega)$ satisfies the Conditions B1-2-3, one can consider the Markov kernel that, for a given pair $(x, u) \in \mathbb{R}^d \times \mathbb{R}^d$, starts by generating $\omega \sim \mathrm{M}(x, \mathrm{d}\omega)$ and then proceeds to Steps 1-2-3 with generalized operator $u \mapsto \widetilde{\mathsf{R}}(u, x, \omega)$. Lemma 1 shows that this leads to a valid algorithm. This in turns shows that, in Section 2.1, one can also consider randomized vector fields: instead of performing a reflection with respect to a fixed vector field $\mathscr{F}(x)$, one can instead generate a random vector $v \in \mathcal{X} \setminus \{0\}$ sampled from a distribution that depends on the vector $x \in \mathcal{X}$ only (i.e. does not depend on the current direction $u \in \mathcal{S}$) and attempts the reflection $u \mapsto \mathscr{R}_v(u)$. This remark also shows the correctness of the methods described in Section 2.4.

A.2 Proof of Proposition 1

Recall that the quantity $\lambda(x, u)$ is defined as $\lambda(x, u) \equiv \langle -\nabla \log \pi(x), u \rangle_+$. Under Assumption (A1) and a discretization parameter $\delta > 0$, the acceptance probability $\alpha^\delta(x, u)$that the proposal $(x, u) \mapsto (x + \delta\, u, u)$ is accepted reads

$$\begin{aligned} \alpha^\delta(x, u) &= \exp\left[\min\left\{0, \log \pi(x + \delta\, u) - \log \pi(x, u)\right\}\right] \\ &= \exp\left\{\delta \times \langle \nabla \log \pi(x), u \rangle_-\right\} \ + \ \mathcal{O}(\delta^2) \\ &= \exp\left\{-\delta \times \lambda(x, u)\right\} \ + \ \mathcal{O}(\delta^2) \end{aligned} \tag{18}$$

where $\mathcal{O}(\delta^2)$ is a quantity whose absolute value is less than a constant times δ^2. For $t > 0$, the probability that the Discrete Bouncy Particle Sampler algorithm accepts $\lfloor t/\delta \rfloor + 1$ consecutive proposals $(x, u) \mapsto (x + \delta\, u, u)$ without reflection attenpts equals $\prod_{k=0}^{\lfloor t/\delta \rfloor} \alpha^\delta(x_k^\delta, u_k^\delta)$. Under Assumption (A3), one can condition upon a fixed trajectory of the Markov process V, i.e. $V_t = v_t$ for all $0 \leq t \leq T$ and $u_k^\delta = \overline{u}_{k\delta}^\delta = v_{k\delta}$, not depending on the parameter δ, so that $\overline{x}_{k\delta}^\delta = \overline{x}_0^\delta + \delta \sum_{j=0}^{k-1} v_{j\delta}$. Equation (18), the continuity of the rate function λ as well as the continuity of the trajectories of the Markov process V, show that

$$\lim_{\delta \to 0} \prod_{k=0}^{\lfloor t/\delta \rfloor} \alpha^\delta(x_k^\delta, u_k^\delta) = \lim_{\delta \to 0} \exp\left\{-\delta \sum_{k=0}^{\lfloor t/\delta \rfloor} \lambda(\overline{x}_{k\delta}^\delta, v_{k\delta})\right\} + \mathcal{O}(\delta) = \exp\left\{-\int_0^t \lambda(\overline{x}_s, v_s)\, \mathrm{d}s\right\}$$

where $\overline{x}_s = x_0 + \int_0^s v_t\, dt$. This means that, in the limit $\delta \to 0$, bounce attempts arrive at rate $\lambda(x, u)$ and, in between the bounces, the limiting process simply evolves according to the dynamics (8).

Finally, once a proposal $(x,u) \mapsto (x+\delta\, u, u)$ is rejected, the second proposal $(x,u) \mapsto (x'',u'')$, i.e. the bounce, is accepted with probability $\alpha_{\text{DR}}(x,u)$ described in Equation (3). By continuity of the density $x \mapsto \pi(x)$, we have that $\pi(x'')/\pi(x) \to 1$ as $\delta \to 0$. Furthermore, the Taylor expansion (18) gives that $1-\alpha^\delta(x,u) = \delta\,\lambda(x,u) + \mathcal{O}(\delta^2)$. Under Assumption, the vector field $\mathscr{F}$ is continuous, which implies that

$$\lim_{\delta\to 0} \frac{1-\alpha^\delta(x'',-u'')}{1-\alpha^\delta(x,u)} \times \frac{\pi(x'')}{\pi(x)} \;=\; \frac{\lambda\{x,-\mathscr{R}_{\mathscr{F}(x)}(u)\}}{\lambda(x,u)}, \tag{19}$$

where we have dropped the dependence on δ from the notation $x'' = x + \delta\, u + \delta\, \mathscr{R}_{\mathscr{F}(x+\delta u)}(u)$ and $u'' = \mathscr{R}_{\mathscr{F}(x+\delta u)}(u)$. It follows from (19) that, in the limit as $\delta \to 0$, a proposed bounce $(x,u) \mapsto (x'',u'')$ is accepted with probability $\mathcal{A}(x,u)$ described in Equation (7), with $(x'',u'') \to (x, \mathscr{R}_{\mathscr{F}(x)}(u))$ as $\delta \to 0$. This completes the proof of Proposition 1.

A.3 Proof of Proposition 2

For a fixed dimension $d \geq 1$, Proposition 1 describes its scaling limit as $\delta \to 0$: the processes $t \mapsto z_t^{d,\delta} = (x_{\delta t}^{d,\delta}, u_{\delta t}^{d,\delta})$ converges on path-space to the jump diffusion that evolves according to $d\overline{X}_t^d = \overline{U}_t^d\, dt$, where $\overline{U}_t^d$ is a Brownian motion on the unit sphere of $\mathbb{R}^d$ whose dynamics is described by the Stochastic Differential Equation (5), in between reflections $\overline{U}_t^d = \mathscr{R}\Big(\overline{U}_{t^-}^d, \nabla \log \pi_d(\overline{X}_{t^-}^d)\Big)$ that arrive at rate $\lambda(\overline{X}_t^d, \overline{U}_t^d) = \Big\langle -\nabla\pi_d(\overline{X}_t^d), \overline{U}_t^d \Big\rangle_+$. The generator of the joint process $(\overline{X}_t^d, \overline{U}_t^d)$ reads

$$\varphi(x,u) \mapsto \langle u, \nabla_x \varphi(x,u)\rangle + \mathcal{L}^{(\kappa,B)}\varphi(x,u) + \lambda(x,u)\,\mathcal{F}\varphi(x,u)$$

where $\mathcal{L}^{(\kappa,B)}$ is the generator of the Brownian motion on the united sphere (5) and $\mathcal{F}$ is the flip operator defined as $\mathcal{F}\varphi(x,u) = \varphi(x,-u) - \varphi(x,u)$. In order to obtain the limit of the process defined in Equation (10), set

$$R_t^d \equiv \|\overline{X}_{d\times t}^d\| - \sigma\, d^{1/2} \qquad \text{and} \qquad \theta_t^d \equiv d^{1/2}\, \Big\langle \overline{X}_{d\times t}^d, \overline{U}_{d\times t}^d \Big\rangle / \|\overline{X}_{d\times t}^d\|.$$

Note that time has been accelerated by a factor d. The process θ_t^d describes the dot product between the position X^d and the direction U^d, scaled by a factor $d^{1/2}$ in order to observe a non-degenerate limiting process. Itô's lemma, neglecting terms of order $1/d$, directly shows (after straightforward algebra) that the Markov process (R_t^d, θ_t^d) has a generator $\mathcal{L}^\varepsilon$ that reads

$$\mathcal{L}^\varepsilon = \varepsilon^{-1}\,\mathcal{L}^{(H)} + \varepsilon^{-1}\,\frac{R}{\sigma^2}\,\theta_+\,\mathcal{F} + \varepsilon^{-2} \underbrace{\left\{\frac{1}{\sigma}\,\mathcal{L}^{(J)} + \frac{\kappa}{2}\,\mathcal{L}^{(K)}\right\}}_{\mathcal{L}^{(\text{Fast})}} \tag{20}$$

with the standard multiscale expansion notation $\varepsilon = 1/\sqrt{d}$, generators $\mathcal{L}^{(J)}$ and $\mathcal{L}^{(K)}$ defined in Equation (11) and

$$\mathcal{L}^{(H)}\varphi = \Big(\theta\,\partial_R - \frac{R}{\sigma^2}\,\partial_\theta\Big)\,\varphi$$

It is important to note that the generator $\mathcal{L}^{(\text{Fast})}$ describes the dynamics of a Markov process that is ergodic with respect to the standard centred Gaussian distribution $\mathrm{G}(\mathsf{d}\theta)$. The dynamics described by the generator (20) is a *fast-slow* systems with slow variable R and fast variable θ. The effective dynamics of the slow variable R as $\varepsilon \to 0$, or equivalently as $d \to \infty$, can be obtained with a standard multiscale expansion (Papanicolaou, 1977; Weinan, 2011), as described for example in Chapter 11 of Pavliotis and Stuart (2008). One seeks a solution $\varphi^\varepsilon(t, R, \theta)$ to the backward Kolmogorov equation $(\partial_t - \mathcal{L}^\varepsilon)\varphi^\varepsilon = 0$ expressed as $\varphi^\varepsilon(t, R, \theta) = \varphi(t, R) + \varepsilon\, A(t, R, \theta) + \varepsilon^2\, B(t, R, \theta) + \mathcal{O}(\varepsilon^3)$ in order to obtain the generator $\mathcal{L}$ describing the leading term φ, i.e. $(\partial_t - \mathcal{L})\varphi = 0$. Expanding the Kolmogorov Equation $(\partial_t - \mathcal{L}^\varepsilon)\varphi^\varepsilon = 0$ in powers of $1/\varepsilon$ shows that

$$\begin{aligned} \mathcal{O}(\varepsilon^{-2}) : &\qquad \mathcal{L}^{(\text{Fast})}\varphi = 0 \\ \mathcal{O}(\varepsilon^{-1}) : &\qquad \mathcal{L}^{(\text{Fast})}A = -\mathcal{L}^{(H)}\varphi \\ \mathcal{O}(1) : &\qquad \partial_t\varphi = \big(\mathcal{L}^{(H)} + \sigma^{-2}\, R\, \theta_+ \mathcal{F}\big)A + \mathcal{L}^{(\text{Fast})}B \end{aligned} \tag{21}$$

Equation $\mathcal{L}^{(\text{Fast})}\varphi = 0$ is immediate since φ does not depend on the variable θ. Furthermore, we have that $\mathcal{L}^{(H)}\varphi = \theta\partial_R\varphi(t, R)$. Consequently, it follows from the second equation of (21) that $A = c(t, R) + g(\theta)\, \partial_R\varphi(t, R)$ where $g : \mathbb{R} \to \mathbb{R}$ is solution of the Poisson equation $\mathcal{L}^{(\text{Fast})}g = -\theta$ and $c(t, R)$ is a function that does not depend upon θ. For an arbitrary function $(t, R, \theta) \mapsto h(t, R, \theta)$, we now use the standard notation $\langle h \rangle_{\mathrm{G}}$ to denote the operation of averaging out the fast variable θ over the standard centred Gaussian distribution $\mathrm{G}(\mathsf{d}\theta) = (2\pi)^{-1/2} e^{-\theta^2/2}\, \mathsf{d}\theta$, i.e. $\langle h \rangle_{\mathrm{G}}(t, R) \equiv \int_{\mathbb{R}} h(t, R, \theta)\, \mathrm{G}(\mathsf{d}\theta)$. We have that $\big\langle \mathcal{L}^{(\text{Fast})}B \big\rangle_{\mathrm{G}} \equiv 0$ so that the second equation in (21) leads to

$$\begin{aligned} \partial_t\varphi(t, R) &= \left\langle \mathcal{L}^{(H)}A + \frac{R}{\sigma^2}\theta_+\mathcal{F}A \right\rangle_{\mathrm{G}} \\ &= \langle \theta\, g(\theta) \rangle_{\mathrm{G}}\, \partial_{RR}\varphi(t, R) + \sigma^{-2}\, R\, \langle -g'(\theta) + \theta_+\, \mathcal{F}g(\theta) \rangle_{\mathrm{G}}\, \partial_R\varphi(t, R) \equiv \mathcal{L}\varphi(t, R). \end{aligned}$$

Algebra and an integration by part (i.e. Stein's Lemma) show that $\langle -g'(\theta) + \theta_+\, \mathcal{F}g(\theta) \rangle_{\mathrm{G}}$ also equals $2\, \langle \theta\, g(\theta) \rangle_{\mathrm{G}}$. Consequently, we have that

$$\partial_t\varphi(t, R) = \langle \theta\, g(\theta) \rangle_{\mathrm{G}} \left(-\frac{R}{\sigma^2/2}\, \partial_R + \partial_{RR} \right)\varphi(t, R) \tag{22}$$

Since g is solution of the Poisson equation $\mathcal{L}^{(\text{Fast})}g(\theta) = -\theta$, the quantity $\langle \theta\, g(\theta) \rangle_{\mathrm{G}}$ also describes the following asymptotic variance,

$$V_\sigma(\kappa) = \lim_{T \to \infty} \mathrm{Var}\left(\frac{1}{\sqrt{2\,T}} \int_0^T \theta_t^\kappa\, \mathsf{d}t \right),$$

where θ_t^κ is the Markov process with generator $\mathcal{L}^{(\text{Fast})} \equiv \frac{1}{\sigma}\, \mathcal{L}^{(J)} + \frac{\kappa}{2}\, \mathcal{L}^{(K)}$. Indeed, Equation (22) is the Kolmogorov backward equation associated to the Ornstein-Uhlenbeck

$$d\overline{R}_t = -2\, \sigma^{-2}\, V_\sigma(\kappa)\, \overline{R}_t\, dt + \sqrt{2\, V_\sigma(\kappa)}\, dW,$$

which concludes the proof of Proposition 2.

B High-dimensional behaviour for finite δ and non-isotropic target

B.1 Proof of Theorem 1

In this section, we use the notation $\Rightarrow$ to denotes convergence in distribution. Since U is a uniformly random unit vector we may write it as $U = Z/\|Z\|$, where Z is a vector of independent standard Gaussians. Further, since $\|Z\|/\sqrt{d} \to 1$ in probability as $d \to \infty$, we henceforth substitute $U = Z/\sqrt{d}$. We also set $\xi_i = \gamma_i X_i$, so the ξ_i are independent and identically distributed with a density of $\exp\{f(\xi_i)\}$. We use the shorthand of $g(x) = f'(x)$ and $h(x) = -f''(x)$, and we let L be the Lipschitz constant for h. Finally we set $\ell(X) \equiv \log \pi(X)$ and define $B(X,U) \equiv \ell(X+\delta U) - \ell(X)$. Firstly,

$$\begin{aligned} B(X,U) &= \delta U^\top \nabla \ell(X) + \frac{\delta^2}{2} U^\top \partial^2 \ell(X) U + \frac{\delta^2}{2} U^\top \left\{ \partial^2 \ell(X+\delta U) - \partial^2 \ell(X) \right\} U \\ &= \frac{\delta}{\sqrt{d}} \sum_{i=1}^d \gamma_i Z_i g(\xi_i) - \frac{\delta^2}{2d} \sum_{i=1}^d \gamma_i^2 Z_i^2 h(\xi_i) + \frac{1}{2\sqrt{d}} T_1^{(d)}, \end{aligned}$$

where $|T_1^{(d)}| \le \delta^3/d \times L \sum_{i=1}^d \gamma_i^3 |Z_i|^3 \to \delta^3 L \times \mathrm{E}[\gamma^3]\,\mathrm{E}[|Z_i|^3] < \infty$. Hence as $d \to \infty$, the Central Limit Theorem gives

$$B(X,U) \Rightarrow \mathsf{N}\left\{ -\frac{1}{2}\delta^2 \mathrm{m}_2(\gamma) J,\ \delta^2 \mathrm{m}_2(\gamma) J \right\}. \tag{23}$$

The result for α_{pu} then follows from Proposition 2.4 of Roberts et al. (1997). The ith component of the gradient vector with respect to which a bounce might occur is

$$\begin{aligned} \partial_{x_i} \ell(X+\delta U) &= \gamma_i g(\gamma_i X_i + \gamma_i \delta U_i) = \gamma_i g(\xi_i + \frac{\delta}{\sqrt{d}} \gamma_i Z_i) \\ &= \gamma_i g(\xi_i) - \frac{\delta}{\sqrt{d}} \gamma_i^2 Z_i h(\xi_i) + R_i^{(d)}, \end{aligned} \tag{24}$$

where $|R_i^{(d)}| \le L\delta^2 \gamma_i^3 Z_i^2/d$. Thus $\|\nabla \ell(X+\delta U)\|^2 = \sum_{i=1}^d \gamma_i^2 g(\xi_i)^2 + \mathcal{O}(1)$, so

$$\|\nabla \ell(X+\delta U)\|^2/d \to \mathrm{E}[\gamma^2 g(\xi)^2] = J\mathrm{m}_2(\gamma).$$

Also, by (24) and the central limit theorem,

$$\begin{aligned} U^\top \nabla \ell(X+\delta U) &= \frac{1}{\sqrt{d}} \sum_{i=1}^d \gamma_i Z_i g(\xi_i) - \frac{\delta}{d} \sum_{i=1}^d \gamma_i^2 Z_i^2 h(\xi_i) + \mathcal{O}(1/\sqrt{d}) \\ &\Rightarrow \mathsf{N}\left\{ -\delta \mathrm{m}_2(\gamma) J, \mathrm{m}_2(\gamma) J \right\}. \end{aligned} \tag{25}$$

Now $U - V = 2[U^\top \nabla \ell(X+\delta U)]/\|\nabla \ell(X+\delta U)\|^2 \times \nabla \ell(X+\delta U)$, so from (25) we have that

$$\begin{aligned} d \times \|U-V\|^2 &= \frac{4d}{\|\nabla \ell(X+\delta U)\|^2} \times [U^\top \nabla \ell(X+\delta U)]^2 \\ &\Rightarrow \frac{4}{J\mathrm{m}_2(\gamma)} [\mathsf{N}\{-\delta \mathrm{m}_2(\gamma) J, \mathrm{m}_2(\gamma) J\}]^2, \end{aligned}$$

so that $\|U - V\| = \mathcal{O}(1/\sqrt{d})$. Consequently, the quantity $\ell(X + \delta U + \delta V) - \ell(X + \delta U)$ equals

$$\delta V^\top \nabla \ell(X + \delta U)) - \frac{1}{2}\delta^2 V^\top \partial^2 \ell(X + \delta U) V + \frac{1}{2\sqrt{d}} T_2^{(d)},$$

where $|T_2^{(d)}| \leq \delta^3 L \sum_{i=1}^d \gamma_i^3 V_i^2 |Z_i| = \mathcal{O}(1)$ by the Lipschitz condition on f'', the boundedness of $\mathrm{E}[\gamma^3]$ and because $\|U - V\|^2 = \mathcal{O}(1/d)$ and $U = Z/\sqrt{d}$. Further, the quantity $-B(X, U)$ also reads

$$\ell(X) - \ell(X + \delta U) = -\delta U^\top \nabla \ell(X + \delta U)) - \frac{1}{2}\delta^2 U^\top \partial^2 \ell(X + \delta U) U + \mathcal{O}(1/\sqrt{d}).$$

Subtracting the two expressions and noting that the bounce is constructed so that $(U+V)^\top \nabla \ell(X + \delta U) = 0$ yields

$$\begin{aligned} \ell(X + \delta U + \delta V) - \ell(X) &= \delta (U + V)^\top \nabla \ell(X + \delta U)) + \frac{1}{2}\delta^2 U^\top \partial^2 \ell(X + \delta U) U \\ &\quad - \frac{1}{2}\delta^2 V^\top \partial^2 \ell(X + \delta U) V + \mathcal{O}(1/\sqrt{d}) \\ &= \frac{1}{2}\delta^2 (U - V)^\top \partial^2 \ell(X + \delta U)(2U - (U - V)) + \mathcal{O}(1/\sqrt{d}). \end{aligned} \tag{26}$$

Since $\|U\| = 1$ and $\|U - V\| = \mathcal{O}(1/\sqrt{d})$, the term to control in (26) is

$$\begin{aligned} \delta^2 |(U - V)^\top \partial^2 \ell(X + \delta U) U| &= 2\delta^2 \frac{U^\top \nabla (X + \delta U)}{\|\nabla \ell(X + \delta U)\|^2} |U^\top \partial^2 \ell(X + \delta U) \nabla \ell(X + \delta U)| \\ &\sim 2\delta^2 \frac{|\mathsf{N}\{-\delta \mathrm{m}_2(\gamma) J, \mathrm{m}_2(\gamma) J\}|}{d J \mathrm{m}_2(\gamma)} |W^{(d)}|, \end{aligned}$$

where

$$\begin{aligned} W^{(d)} &\equiv \frac{1}{\sqrt{d}} \sum_{i=1}^d Z_i \gamma_i^3 h(\xi_i + \gamma_i Z_i/\sqrt{d}) \gamma_i g(\xi_i + \gamma_i Z_i/\sqrt{d}) \\ &= \frac{1}{\sqrt{d}} \sum_{i=1}^d \{Z_i \gamma_i^3 h(\xi_i) \gamma_i g(\xi_i) + \mathcal{O}(1/\sqrt{d})\}. \end{aligned}$$

But $\mathrm{E}\{\sum_{i=1}^d Z_i \gamma_i^3 h(\xi_i) g(\xi_i)\} = 0$ so $W^{(d)} = \mathcal{O}(1)$, and, hence, $\ell(X + \delta U + \delta V) - \ell(X) = \mathcal{O}(1/\sqrt{d})$. It follows that

$$\ell(X + \delta U) - \ell(X + \delta U + \delta V) = \ell(X + \delta U) - \ell(X) - \{\ell(X + \delta U + \delta V) - \ell(X)\} \to B(X, U)$$

in probability. If there is a delayed-rejection event then the standard move must have been rejected and, for example, B must be negative. Let DR be the event that the standard move has been rejected and so a delayed-rejection step is being attempted. Let $f_B(b)$ be the *a priori* density for B at stationarity, and let $f_{B|\mathsf{DR}}(b)$ be the density conditional on there being a delayed-rejection event. Then

$$f_{B|\mathsf{DR}}(b) = \frac{f_B(b)\,\{1 - 1 \wedge \exp(b)\}}{\int_{-\infty}^{\infty} f_B(b)\,\{1 - 1 \wedge \exp(b)\}\,\mathrm{d}b}, \tag{27}$$

which is well-behaved and has no mass where $\alpha_{\text{dr}}(X^{(d)}, U^{(d)})$ is undefined. In the limit as $d \to \infty$, $f_B(b)$ is the density of the Gaussian distribution in (23), and (27) gives the limiting conditional density. It follows from the Bounded Convergence Theorem that

$$\alpha_{\text{dr}}(X, U) = \mathrm{E}\left[1 \wedge \frac{1 - 1 \wedge \exp\{\ell(X + \delta U) - \ell(X + \delta U + \delta V)\}}{1 - 1 \wedge \exp\{\ell(X + \delta U) - \ell(X)\}} \exp\{\ell(X + \delta U + \delta V) - \ell(X)\} \mid \mathsf{DR}\right]$$
$$\to \mathrm{E}\left[1 \wedge \frac{1 - \{1 \wedge \exp(B)\}}{1 - \{1 \wedge \exp(B)\}} \mid \mathsf{DR}\right] = 1,$$

as required.

B.2 Simulation study varying d for fixed δ

In dimension d we explore a $\mathsf{N}\{0, \mathrm{diag}(\sigma_1^2, \ldots, \sigma_d^2)\}$ target with $\gamma_i^2 = 1/\sigma_i^2 = 2i/(d+1)$. In this case, $J = 1$ and, essentially, $\mathrm{m}_2(\gamma) = 1$, so that for fixed δ, Theorem 1 tells us that, asymptotically, we expect $\alpha_{\text{pu}} \to 2\Phi(-\delta/2)$ and $\alpha_{\text{dr}} \to 1$. For each combination of $d \in \{5, 10, 20, 50, 100, 200\}$ and $\delta \in \{0.1, 0.2, 0.5, 1.0, 2.0, 3.0\}$ (except $d = 5$ and $\delta = 0.1$) where the Monte Carlo relative error was large) three replicate runs were performed with $\kappa = 1.0$. Empirical acceptance rates for the standard moves and for the delayed-rejection move were recorded for each run.

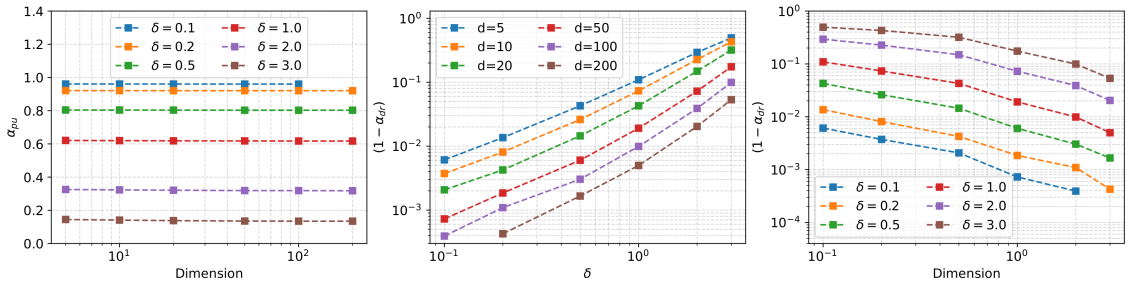

Figure 4: Left: α_{pu} against $\log_{10} d$ split by δ; centre: $\log_{10}(1 - \alpha_{\text{dr}})$ against $\log_{10} d$ split by δ; right: $\log_{10}(1 - \alpha_{\text{dr}})$ against $\log_{10} \delta$ split by d.

The left panel of Figure 4 plots α_{pu} against the dimension d for each fixed δ and demonstrates that for each fixed δ, the acceptance rate α_{pu} is insensitive to d, suggesting that the asymptotics are highly relevant even for very low dimension. The central and right panels plot $(1 - \alpha_{\text{dr}})$ against the dimension d and $(1 - \alpha_{\text{dr}})$ against δ respectively, and suggests that, not only does $\alpha_{\text{dr}} \to 1$ as $d \to \infty$ but that, at least in this example, asymptotically, $1 - \alpha_{\text{dr}} \propto \delta^2/d$.

C Further simulation studies

C.1 Convergence from the tails

We illustrate the contrast between the discrete bouncy particle sampler and Hamiltonian Monte Carlo on a fifty-dimensional target with a density described in Equation (16). We first tuned each algorithm by starting at random positions of the form $(\sigma_1 z_1, \ldots, \sigma_d z_d) \times r_\star$, where $z \in \mathbb{S}^{d-1}$ is a

uniformly random unit vector in $\mathbb{R}^d$. For Hamiltonian Monte Carlo we tried different integration times, T, and for each T, we followed the advice of Beskos et al. (2013), noting that convergence to the optimal acceptance rate as dimension increases is slow, and tuned the number of leapfrog steps so as to achieve an acceptance probability of a little over 70%. For the discrete bouncy particle sampler we tried different values for δ, and for each δ we chose κ so that, as suggested in Section 4.1, the mean dot product diagnostic is around 0.3-0.4. This suggested that $(T = 2.0, L = 4)$ for Hamiltonian Monte Carlo and $(\delta = 2.0, \kappa = 0.7)$ for the discrete bouncy particle sampler were reasonable tunings. Since, for a non-toy target in $d = 50$, gradient evaluations will be much more costly than likelihood evaluations it is reasonable to a first approximation to assess efficiency by comparing the ratio of effective sample size to the average number of gradient evaluations per iteration. With 10^5 iterations this is ≈ 1473 for Hamiltonian Monte Carlo and ≈ 1014 for the discrete bouncy particle sampler. However, the apparent relative success of Hamiltonian Monte Carlo disguises a serious underlying issue.

We reran Hamiltonian Monte Carlo for 10^6 iterations 40 additional times with $X_0 = \gamma(\sigma_1 z_1, \ldots, \sigma_d z_d) \times r_\star$ for each $\gamma \in \{1.5, 2.0.2.5, 3.0\}$, with a new, independent z vector on each of the 160 occasions. On each occasion we counted the fraction of times where, by iteration 10^6 the algorithm had ever had a value with $\|x\|_{\mathrm{M}} \leq r_\star$; i.e., the algorithm had reached the main posterior mass. The number of runs which converged by this measure were: $40/40$ $(\gamma = 1.5)$, $36/40$ $(\gamma = 2.0)$, $4/40$ $(\gamma = 2.5)$ and $0/40$ $(\gamma = 3.0)$; indeed, for every run with $\gamma = 3.0$ the empirical acceptance rate was exactly zero. This fits with the known lack of geometric ergodicity of Hamiltonian Monte Carlo on light-tailed targets. By contrast, for the discrete bouncy particle sampler with $\gamma = 3.0$, all 40 runs converged within 1000 iterations, and, indeed, 26 of the runs converged within 300 iterations.

In summary, on this occasion, when both algorithms were started from the main posterior mass, the discrete bouncy particle sampler was competetive with Hamiltonian Monte Carlo, though less efficient. However, because our algorithm depends on $\nabla \log \pi$ only through the unit vector, it is robust to large $\|\nabla \log \pi\|$, unlike Hamiltonian Monte Carlo.

C.2 The Markov modulated Poisson process

Finally, we consider a k-state, continuous-time Markov chain Z_t started from state 1, and a Poisson process N_t whose rate λ_t is a fixed function of Z_t. The doubly-stochastic process is parameterized by the rate matrix for the Markov chain, Q, and a vector of rates for the Poisson process, λ, where $\lambda_i,\ (i = 1, \ldots, k)$ is the rate of N_t when $Z_t = i$.

The event times of N_t are observed over a time window $[0, t_{\mathrm{end}}]$, but the behaviour of Z_t is unknown, and we wish to perform inference on (Q, λ). Setting $\Lambda = \mathrm{diag}(\lambda)$, the likelihood for the number of events n and the event times $t_1, \ldots, t_n$ is (Fearnhead and Sherlock, 2006):

$$L(Q, \lambda; t) = \mathrm{e}' \exp[(Q - \Lambda)t_1]\Lambda \exp[(Q - \Lambda)(t_2 - t_1)]\Lambda \ldots \Lambda \exp[(Q - \Lambda)(t_{\mathrm{end}} - t_n)]1,$$

where 1 is the k-vector of ones and $\mathrm{e}' = (1, 0, \ldots, 0)$. We simulated a dataset using a cyclic four-state Markov chain for a 200-second time window with Q parameters of: $Q_{12} = Q_{23} = Q_{41} = 1.0,\ Q_{34} = 0.25$ and all other off-diagonal rates set to zero. The rate parameters were $\lambda_1 = 20.0, \lambda_2 = 5.0,\ \lambda_3 = 1.0$ and $\lambda_4 = 10.0$. We then conducted inference on the natural logarithm of each parameter that was not systematically zero, placing independent $N(0, 2^2)$ priors on each of these.

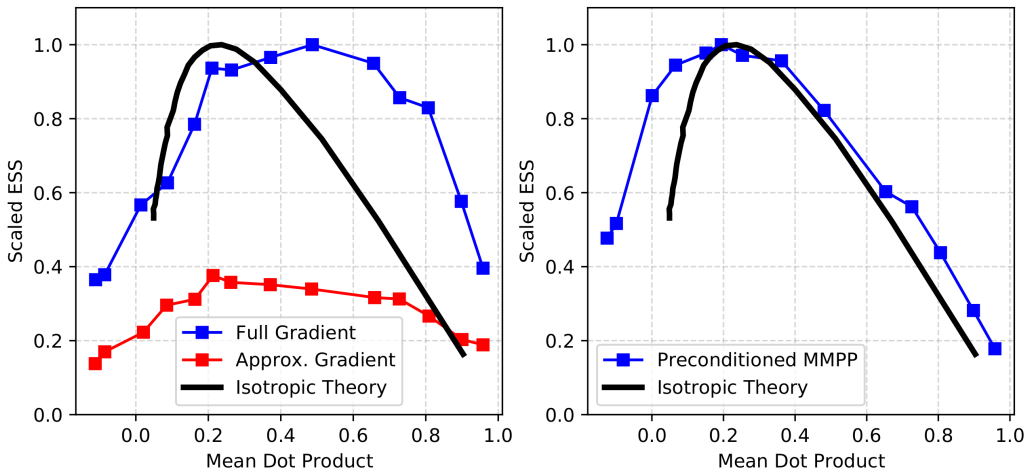

Figure 5: Scaled effective sample size as a function of the mean dot product for several values of κ, obtained from Discrete Bouncy Particle Sampler runs of length $N = 2 \times 10^5$ applied to the logarithms of the MMPP parameters with no preconditioning and with $n_{\text{cpt}} = d = 8$ and $n_{\text{cpt}} = 3$ (Left), and with crude preconditioning and $n_{\text{cpt}} = d$ (Right).

Code was written in C++ where auto-differentiation was not available for general matrix exponentials, and so numerical differentiation via centred differences was used (the cheaper, first-order Euler approximation led to precision problems). We applied the Discrete Bouncy Particle Sampler for 2×10^5 iterations for a number of κ values and repeated this but evaluating only $n_{\text{cpt}} = 3$ randomly-orientated components of the eight-dimensional gradient vector on each delayed-rejection step. Figure 5 plots scaled effective sample size against κ and suggests that the optimal mean dot product is around 0.5 when all gradient components are used and around 0.2 when three random components are used. The optimal effective sample size in the latter case is around $3/8$ of the former; since the number of gradient calculations during a bounce is also reduced by $3/8$ this suggests no loss in overall efficiency. When all components are used, tuning to a dot product of 0.2 brings only a 10% reduction in effective sample size. The posterior variance matrix has a condition number of 49.2, so, following typical practice, for each κ the Discrete Bouncy Particle Sampler was rerun using a crude preconditioning matrix (Section 2.5) of $M = \text{diag}(1/2, 2, 1, 1, 2, 2, 2, 2)$. Although the effective variance matrix still has a condition number of ≈ 4.3 the right panel of Figure 5 shows that the optimal choice of κ is now ≈ 0.2.