


TECHNICAL REPORTS FROM THE ELECTRONICS GROUP
AT THE UNIVERSITY OF OTAGO

**Metropolis-Hastings algorithms with autoregressive proposals,
and a few examples**

by

Richard A. Norton, Colin Fox

ELECTRONICS TECHNICAL
REPORT No. 2016-1



UNIVERSITY OF OTAGO
DUNEDIN, NEW ZEALAND

Electronics Group at Otago

In 1987 Millman and Grabel discarded the historical definition of ‘electronics’ as the science and technology of the motion of charges, preferring instead the operational definition that the primary concern of people doing electronics is *information processing*. This makes a distinction from *energy processing* practiced in the rest of electrical engineering. The act of information processing is what gets electronics practitioners involved in the fours ‘C’s: communication, computation, control, and components. This practical definition seems to describe well the activities within the Electronics Group in the Physics Department at the University of Otago, and the range of topics covered in this technical report series.

In May 2016, research within the Electronics Group include projects applying inference algorithms to embedded sensors, on lightweight GPS tags for birds, modelling and control of a robotic elbow, design and deployment of an under-sea glider, analysis of networks of random resistors, electrical impedance imaging, calibration of numerical models for geothermal fields using Bayesian inference, modelling and sampling of Gaussian processes, and efficient algorithms for Markov chain Monte Carlo applied to inverse problems.

Citing this Report

This report should be cited as:

Richard A. Norton and Colin Fox, “Metropolis-Hastings algorithms with autoregressive proposals, and a few examples”, *Electronics Technical Reports* No. 2016-1, ISSN 1172-496X, May 2016.

Alternatively, using BibTeX, as:

```
@techreport{etr2016-1,  
  title{{Metropolis-Hastings algorithms with autoregressive  
  proposals, and a few examples}},  
  author={Richard A. Norton and Colin Fox},  
  series={Technical Reports from the Electronics Group at the  
  University of Otago},  
  institution={{University of Otago}},  
  issn={1172-496X (Print) 1172-4234 (Online)},  
  url={http://www.physics.otago.ac.nz/reports/electronics/ETR2016-1.pdf},  
  number={2016-1},  
  date{May 2016} }
```

Metropolis-Hastings algorithms with autoregressive proposals, and a few examples

Richard A. Norton & Colin Fox

Abstract

We analyse computational efficiency of Metropolis-Hastings algorithms with stochastic AR(1) process proposals. These proposals include, as a subclass, discretized Langevin diffusion (e.g. MALA) and discretized Hamiltonian dynamics (e.g. HMC).

We derive expressions for the expected acceptance rate and expected jump size for MCMC methods with general stochastic AR(1) process proposals for the case where the target distribution is absolutely continuous with respect to a Gaussian and the covariance of the Gaussian is allowed to have off-diagonal terms. This allows us to extend what is known about several MCMC methods as well as determining the efficiency of new MCMC methods of this type. In the special case of Hybrid Monte Carlo, we can determine the optimal integration time and the effect of the choice of mass matrix.

By including the effect of Metropolis-Hastings we also extend results by Fox and Parker, who used matrix splitting techniques to analyse the performance and improve efficiency of stochastic AR(1) processes for sampling from Gaussian distributions.

Contents

1	Introduction	5
2	Preliminary results and notation	11
2.1	Stochastic AR(1) processes correspond to matrix splittings	11
2.2	Notation	12
3	Gaussian targets	13
3.1	Expected acceptance rate for a Gaussian target	13
3.2	Expected squared jump size for a Gaussian target	16
4	Non-Gaussian targets	19
4.1	Expected acceptance rate for a non-Gaussian target	19
4.2	Expected jumpsize for non-Gaussian target	25
5	Examples	27
5.1	Discretized Langevin diffusion - MALA and SLA	27
5.2	Discretized Langevin diffusion - more general algorithms	29
5.3	L -step methods	31
5.4	Hybrid Monte Carlo	34
6	Concluding remarks	39
 Appendix		
A	Proofs	41
A.1	Proof of Lemma 3.1.1	41
A.2	Proof of Theorem 5.2.2	41
A.3	Proof of Lemma 5.3.1	43

A.4	Proof of Theorem 5.3.1	44
A.5	Proof of Theorem 5.4.1	45
A.6	Proof of Corollary 5.4.1	45
A.7	Proof of Theorem 5.4.2	46
A.8	Proof of Theorem 5.4.4	47

Chapter 1

Introduction

We consider Metropolis-Hastings (MH) algorithms for sampling from a target distribution π_d using a stochastic AR(1) process proposal; given current state $x \in \mathbb{R}^d$ the proposal $y \in \mathbb{R}^d$ is given by

$$y = Gx + g + \nu \quad (1.1)$$

where $G \in \mathbb{R}^{d \times d}$ is the iteration matrix, $g \in \mathbb{R}^d$ is a fixed vector and ν is an independent and identically distributed (i.i.d.) draw from $N(0, \Sigma)$. In general, G , g and Σ may depend on x . We will refer to (1.1) as an *AR(1) proposal* or *stochastic AR(1) proposal*. The proposal is accepted with probability

$$\alpha(x, y) = 1 \wedge \frac{\pi_d(y)q(y, x)}{\pi_d(x)q(x, y)}$$

where $\pi_d(x)$ denotes the target probability density function, $q(x, dy) = q(x, y)dy$ is the transition kernel for the proposal y given current state x , and $p \wedge q = \min\{p, q\}$.

Algorithms using AR(1) proposals include: the random-walk Metropolis algorithm (RWM) [29], the simplified Langevin algorithm (SLA) [9], the so-called θ -SLA method [9], preconditioned versions of RWM and SLA [9], and the Crank-Nicolson (CN) and preconditioned Crank-Nicolson (pCN) proposals [15]. When the target distribution is Gaussian, then the Metropolis-adjusted Langevin algorithm (MALA) [30] and the Hybrid Monte Carlo algorithm (HMC) [16, 6, 26] can also be written in the form of (1.1), and MALA is the same as SLA. For any target distribution, one step of HMC is the same as MALA [6].

Analysis of these Markov chain Monte Carlo (MCMC) methods is almost exclusively limited to the case when the target distribution π_d is a change of measure from a reference product distribution so that

$$\frac{d\pi_d}{d\tilde{\pi}_d}(x) = \exp(-\phi_d(x)) \quad (1.2)$$

for some $\phi_d : \mathbb{R}^d \mapsto \mathbb{R}$ where $\tilde{\pi}_d$ is a product distribution of the form

$$\tilde{\pi}_d(x) = \prod_{i=1}^d \lambda_i f(\lambda_i x_i) \quad (1.3)$$

for some $f : \mathbb{R} \mapsto \mathbb{R}$ and sequence $\{\lambda_i\}_{i=1}^d \subset \mathbb{R}$. We can view π_d and $\tilde{\pi}_d$ as finite-dimensional approximations of infinite-dimensional measures π and $\tilde{\pi}$ satisfying $\frac{d\pi}{d\tilde{\pi}}(x) = \exp(-\phi(x))$ on some state space.

Examples of inverse problems that yield posterior distributions of this form can be found in [34, 9, 8].

Analysis of MH algorithms with AR(1) proposals is also typically limited to the cases when $d = 1$ or $d \rightarrow \infty$. Of which the $d \rightarrow \infty$ case is more important because it is used as an approximation for the practical computational problem when d is large but finite. Precisely what is meant by ‘large’ is problem dependent. For example, [30, §3] demonstrates that for MALA with $\phi_d = 0$, $f(x) \propto \exp(-x^2/2)$ and $\lambda_i = 1$, then $d = 5$ is large, but if f is non-symmetric then large d is greater than 10.

Analyses of RWM and MALA began with the case when $\phi = 0$ and $\lambda_i = 1$ for all i (so that π_d has product form with i.i.d. components) based on discretizations of a Langevin diffusion process. Roberts, Gelman and Gilks [29] for RWM and then Roberts and Rosenthal [30] for MALA showed that as $d \rightarrow \infty$, the first component of the Markov chain converges to a Langevin diffusion process and the ‘speed’ of the diffusion process is maximised when the acceptance rate is 0.234 for RWM and 0.574 for MALA. This is equivalent to maximising the expected squared jump size of the Markov chain. The non i.i.d. cases when $\phi = 0$ and $\lambda_i \neq 1$ for RWM and MALA are subsequently treated in [31, 3, 4, 5]. In these articles it is noted that while the optimal acceptance rate for RWM and MALA remains the same, the expected jump size of RWM decreases as the l^2 -norm of the sequence $\{\lambda_i\}$ increases, while MALA depends on the l^6 -norm. An ‘inhomogeneous’ RWM proposal is also considered in [3], which is what we will call ‘preconditioning’. In [5], it is also noted that their results also hold when the target distribution is a multivariate normal whatever the covariance matrix since orthogonal transformations can transform the target to one with independent components. We will exploit this fact throughout this article. A non-product form of target for RWM is considered in [14].

The case of RWM and SLA (a simplified version of MALA) for non-product target distributions, when $\phi \neq 0$, is considered in [9]. The optimal acceptance rates remain 0.234 for RWM and 0.574 for SLA (same as MALA), under certain conditions on ϕ_d and $\{\lambda_i\}_{i=1}^\infty$. Again, the expected jump sizes for these algorithms decrease with the l^2 - and l^6 -norms of $\{\lambda_i\}$. It is also suggested how to precondition the RWM and SLA proposals in [9].

HMC was analysed in [6] for the case when the target distribution has product form with i.i.d. components. Similarly to the analyses of RWM, MALA and SLA, the authors of [6] showed that the expected squared jump size for HMC is maximised when the acceptance rate is 0.651, and this corresponds to $\mathcal{O}(d^{1/4})$ steps to traverse state space. This compares favourably with RWM and MALA which require $\mathcal{O}(d)$ and $\mathcal{O}(d^{1/3})$ steps respectively for the same problem, but still blows up as $d \rightarrow \infty$.

More recent analyses has shown that some methods can be modified so they are well-defined in the infinite-dimensional function space setting, so in the limit as $d \rightarrow \infty$, the methods achieve a positive acceptance rate without zero step size and only $\mathcal{O}(1)$ steps are required to traverse state space. This requires modifying the proposals by ‘preconditioning’ and/or a coordinate transformation. For example, in the case when the target is a change of measure from a Gaussian reference measure, when

$$\tilde{\pi}_d(x) \propto \exp\left(-\frac{1}{2}x^T A x + b^T x\right), \quad (1.4)$$

for symmetric positive definite matrix $A \in \mathbb{R}^{d \times d}$ and vector $b \in \mathbb{R}^d$, the CN and pCN proposals are analysed in [15], and a variant of HMC is analysed in [7]. Other examples of this approach include [8, 10, 28]. Another way to view preconditioning of MALA and HMC is given by [20].

In pCN and the variant of HMC in [7] it is necessary to draw independent samples from $N(0, A^{-1})$ or compute a spectral decomposition of A which effectively transforms the reference

measure to a product distribution; both of which could be computationally infeasible when d is large. In CN, the action of $(I + tA)^{-1}$ for some $t > 0$ is required per iteration of the Markov chain, see [15], which may be expensive to compute.

The results in [9], where $\tilde{\pi}_d$ has product form, easily extend to the case where $\tilde{\pi}_d$ is Gaussian (1.4) where A may have non-zero off-diagonal terms. This is obvious once we recognise that the Markov chains for RWM and SLA are invariant to orthogonal coordinate transformations. That is, there exists an orthogonal coordinate transformation that diagonalizes the covariance matrix of π_d and G and Σ in (1.1), see e.g. Lemma 3.1.2. It is important to note that it is not necessary to compute the orthogonal transformation, as it is enough to simply know that it exists, and the efficiency of the untransformed Markov chain is identical to the transformed chain.

We will extend this idea to MH algorithms with general AR(1) proposals where G and Σ are functions of A , targeting distributions that are either Gaussian, or a change of measure from a reference Gaussian distribution (1.4). In particular, the Gaussian reference distribution is allowed to have off-diagonal terms so it is not restricted to product form.

Therefore, we extend the study of MH algorithms with particular AR(1) proposals to general AR(1) proposals where G and Σ are functions of the reference precision matrix A . We also extend the study of MH algorithms with AR(1) proposals targeting distributions defined by (1.2) and (1.3) to target distributions defined by (1.2) and (1.4) where A may have off-diagonal terms.

Another important feature of this analysis is that the proposals do not necessarily require independent samples from $N(0, A^{-1})$ or $N(0, A)$, multiplying by $A^{1/2}$ or $A^{-1/2}$, or computing a spectral decomposition of A or A^{-1} ; even though we use the existence of a spectral decomposition of A for theoretical purposes. This fact separates this new theory from previous theory for CN, pCN and HMC in [15, 7], where these proposals include operations that may be computationally infeasible in high dimensions.

By generalising the results in [9, 6] under some assumptions, we calculate limits for the expected acceptance rate and expected squared jump size for MH algorithms with AR(1) proposals as $d \rightarrow \infty$. We can then decide on the efficiency of a method based on expected jump size and the computing cost for each proposal.

Our new theory encompasses existing MCMC methods with AR(1) proposals, which are now special cases for our theory, and we can extend the results that are currently available for SLA, HMC, θ -SLA, and preconditioned versions of these MCMC methods, see Section 5. In the case of HMC for a Gaussian target, we are no longer restricted, as in [6], to an i.i.d. product target, and we now have criteria for how to choose the mass matrix (preconditioner) and the total time to integrate the Hamiltonian system. Previous analyses of HMC only provided guidance on tuning the time step until the acceptance rate is 0.651.

We can also apply our new theory to new MCMC methods. For example we can analyse an MCMC method where the proposal is L steps of the SLA proposal before accepting or rejecting. We show that for any L , the step size should be tuned until the acceptance rate is 0.574, the same as MALA and SLA, and when the computing cost is dominated by matrix-vector products with A , then it is optimal to use $L > 1$. Moreover, as the cost of evaluating ϕ_d increases, so does the optimal choice of L .

Our analysis relies on the theory of *matrix splitting* which originated in numerical linear algebra for iteratively solving linear systems of equations [2], but has since been applied to

sampling from Gaussian distributions [17, 19, 18]. As we will see in Section 2.1, if the spectral radius of G is less than 1, then it is possible to rewrite (1.1) in terms of a matrix splitting of a matrix \mathcal{A} , which is not equal to A in general. By defining *splitting* matrices M and N such that $\mathcal{A} = M - N$ then y from (1.1) satisfies

$$My = Nx + \beta + \nu \tag{1.5}$$

where β is a vector, ν is an i.i.d. draw from $N(0, M^T + N)$, $G = M^{-1}N$, $g = M^{-1}\beta$ and $\Sigma = M^{-1}(M^T + N)M^{-T}$. The converse statement, y satisfies (1.1) if y satisfies (1.5), only requires that M^{-1} exists. Moreover, if the spectral radius of G is less than 1, then the Markov chain generated by (1.1) or (1.5) without the MH accept/reject step, will converge to $N(\mathcal{A}^{-1}\beta, \mathcal{A}^{-1})$, which we call the *proposal limit distribution*, see [19]. We call this Markov chain the *proposal chain*.

Fox and Parker [19] realised that Gibbs sampling from a Gaussian $N(\mathcal{A}^{-1}\beta, \mathcal{A}^{-1})$ is very closely related to the Gauss-Seidel iterative solution to a linear system of equations $\mathcal{A}x = \beta$, and that Gauss-Seidel and Gibbs sampling use the same matrix splitting $\mathcal{A} = M - N$; $M = L + D$ and $N = U$ where L , D and U are the strictly lower triangular, diagonal, and upper triangular parts of \mathcal{A} respectively. A generalisation of this observation is that all proposal chains generated by (1.5) are generalised fixed-scan Gibbs samplers for a Gaussian.

If (1.1) or (1.5) is a proposal for the MH algorithm then the transition kernel changes from that of the proposal chain and we cannot use the theory in [19] to determine its convergence properties. Moreover, acceleration techniques suggested in [18, 19] for the proposal chain may not accelerate the MH algorithm. For example, Goodman and Sokal [21] accelerated Gibbs sampling of normal distributions using ideas from multigrid linear solvers, but only observed modest efficiency gains in their non-normal examples (exponential distributions with fourth moments). Also, Green and Han [22] applied successive-over-relaxation to a local Gaussian approximation of a non-Gaussian target distribution as a proposal for the MH algorithm. Again, they did not observe significant acceleration in the non-normal target case.

This article is useful for designing efficient MCMC methods because it shows how the expected squared jump size depends on the eigenvalues of G and the difference between the proposal limit $N(\mathcal{A}^{-1}\beta, \mathcal{A}^{-1})$ and the target reference $N(A^{-1}b, A^{-1})$. Generally, an efficient MH algorithm with an AR(1) proposal should satisfy:

1. $\mathcal{A} = A$ and $\beta = b$ or have small differences.
2. The spectral radius of G should be bounded well below 1.
3. The proposal should be cheap to compute, i.e. multiplying by G or M^{-1} and independent sampling from $N(0, \Sigma)$ or $N(0, M^T + N)$ should be cheap to compute.

In addition to providing these ‘rules of thumb’, we have quantified the effect of the eigenvalues of G , the difference between the eigenvalues of A and \mathcal{A} , and the difference between $A^{-1}b$ and $\mathcal{A}^{-1}\beta$.

Our long-term goal is to construct an AR(1) proposal for the MH algorithm based on local Gaussian approximations to the target π_d , using ideas in this article. For example, we might choose A and b such that $-\frac{1}{2}x^T \mathcal{A}x + \beta^T x$ is a local quadratic approximation to $\log \pi_d$, and then choose M and N to define a proposal such that $\mathcal{A} = A$ and $\beta = b$ and the proposal is cheap to

compute. To a certain extent, choosing A and b to obtain a local quadratic approximation to $\log \pi_d$ mimics the design of some optimization algorithms, see e.g. [27], so optimization theory could be a source of inspiration for designing sampling algorithms.

Our analysis is limited to cases where M and N are functions of A . This allows us to simultaneously diagonalise both the AR(1) proposal and the target reference distribution with a coordinate transformation and define a parallel Markov chain (that we never compute with) such that it has the same convergence properties as the original. This is not an overly restrictive condition if the dimension is high enough to make a spectral decomposition of A impractical to compute. We will see below that it includes several important examples of MH algorithms already in use.

The remaining sections are as follows. In Section 2.1 we show that (1.1) and (1.5) are equivalent, then Section 3 presents new analyses for the expected acceptance rate and jump size of MH algorithms with AR(1) proposals when the target distribution is Gaussian ($\phi_d = 0$). We then extend these results to the non-Gaussian case in Section 4. Section 5 then applies this new analysis to proposals from Langevin diffusion and Hamiltonian dynamics. We see that these proposals are AR(1) proposals and we identify the corresponding matrix splitting and proposal limit distribution. Using our earlier analysis we assess the convergence properties of these methods as $d \rightarrow \infty$. We provide concluding remarks in Section 6. Several proofs have been moved to the Appendix to improve readability.

Chapter 2

Preliminary results and notation

2.1 Stochastic AR(1) processes correspond to matrix splittings

We can express a stochastic AR(1) process using either (1.1) or (1.5), provided it converges.

Theorem 2.1.1 If we are given G , g and Σ , and the spectral radius of G is less than 1, then the stochastic AR(1) process (1.1) can be written as (1.5) using

$$\mathcal{A} = \left(\sum_{l=0}^{\infty} G^l \Sigma (G^T)^l \right)^{-1} \quad (2.1)$$

and

$$\begin{aligned} M &= \mathcal{A}(I - G)^{-1}, \\ N &= \mathcal{A}(I - G)^{-1}G, \quad \beta = \mathcal{A}(I - G)^{-1}g. \end{aligned} \quad (2.2)$$

Note that $\mathcal{A} = M - N$ is symmetric and positive definite (spd).

Proof. Since the spectral radius of G is less than 1 and Σ is spd, it follows that $\mathcal{A}^{-1} := \sum_{l=0}^{\infty} G^l \Sigma (G^T)^l$ is well-defined and spd. Then (2.1) and (2.2) satisfy $\mathcal{A} = M - N$, $G = M^{-1}N$ and $g = M^{-1}\beta$. We must also check that $\Sigma = M^{-1}(M^T + N)M^{-T}$ is satisfied. Substituting (2.2) into $M^{-1}(M^T + N)M^{-T}$ we get

$$\begin{aligned} M^{-1}(M^T + N)M^{-T} &= (I - G)\mathcal{A}^{-1} + G\mathcal{A}^{-1}G^T \\ &= \mathcal{A}^{-1} - G\mathcal{A}^{-1}G^T \\ &= \sum_{l=0}^{\infty} G^l \Sigma (G^T)^l - \sum_{l=1}^{\infty} G^l \Sigma (G^T)^l \\ &= \Sigma. \end{aligned}$$

■

If we are given M , N and β , and M^{-1} exists, then it is obvious that if y satisfies (1.5), then y satisfies (1.1) with $G = M^{-1}N$, $g = M^{-1}\beta$ and $\Sigma = M^{-1}(M^T + N)M^{-T}$.

We also remark that Theorem 2.1.1 does not apply to RWM since $G = I$ for RWM.

In the following special case we obtain a symmetric matrix splitting.

Corollary 2.1.1 If the spectral radius of G is less than 1 and $G\Sigma$ is symmetric, then the stochastic AR(1) process (1.1) has a corresponding matrix splitting defined by

$$\begin{aligned} M &= \Sigma^{-1}(I + G), & \mathcal{A} &= M(I - G) = \Sigma^{-1}(I - G^2), \\ N &= MG = \Sigma^{-1}(I + G)G, & \beta &= Mg = \Sigma^{-1}(I + G)g, \end{aligned}$$

and M and N are symmetric (we say the matrix splitting is symmetric).

Proof. These matrix splitting formulae follow from the identity

$$\sum_{l=0}^{\infty} G^l \Sigma (G^T)^l = \sum_{l=0}^{\infty} G^{2l} \Sigma = (I - G^2)^{-1} \Sigma.$$

To see that M is symmetric (and hence also N since \mathcal{A} is symmetric and $\mathcal{A} = M - N$) we note that

$$M = \left((I - G) \sum_{l=0}^{\infty} G^{2l} \Sigma \right)^{-1}.$$

■

2.2 Notation

Throughout the remainder of this article we will use the following notation. Let G_i , λ_i^2 and $\tilde{\lambda}_i^2$ be eigenvalues of G , A and \mathcal{A} respectively. Also define

$$\begin{aligned} \tilde{g}_i &:= 1 - G_i, & g_i &:= 1 - G_i^2, & m_i &:= (A^{-1}b)_i, & \tilde{m}_i &:= (\mathcal{A}^{-1}\beta)_i, \\ r_i &:= \frac{\lambda_i^2 - \tilde{\lambda}_i^2}{\lambda_i^2}, & \tilde{r}_i &:= \frac{\lambda_i^2}{\tilde{\lambda}_i^2}, & \hat{r}_i &:= m_i - \tilde{m}_i \end{aligned}$$

and

$$\begin{aligned} T_{0i} &:= \hat{r}_i^2 \lambda_i^2 (\tfrac{1}{2} r_i g_i - \tilde{g}_i), & T_{1i} &:= \hat{r}_i \lambda_i (r_i g_i - \tilde{g}_i), & T_{2i} &:= \hat{r}_i \lambda_i (\tilde{r}_i g_i)^{1/2} (1 - r_i G_i), \\ T_{3i} &:= \tfrac{1}{2} r_i g_i, & T_{4i} &:= -\tfrac{1}{2} r_i \tilde{r}_i g_i, & T_{5i} &:= -r_i G_i (\tilde{r}_i g_i)^{1/2}. \end{aligned}$$

In general, all of these quantities may depend on d . The standard normal cumulative distribution function will always be Φ .

We will say that $f_{d,i} = \mathcal{O}(g_{d,i})$ (uniformly in i) as $d \rightarrow \infty$ if for all i and all sufficiently large d , $f_{d,i}/g_{d,i}$ is bounded by a constant that is independent of d and i . Likewise, $f_{d,i} = o(g_{d,i})$ (uniformly in i) as $d \rightarrow \infty$ if $\max_{1 \leq i \leq d} f_{d,i}/g_{d,i} \rightarrow 0$ as $d \rightarrow \infty$. For brevity we will sometimes omit “uniformly in i ”.

Other articles use λ_i^2 to denote the eigenvalues of the covariance matrix corresponding to $\tilde{\pi}_d$ [9, 11, 10, 7]. We do not follow this convention and instead use λ_i^2 to denote eigenvalues of the precision matrix. Since sampling from the Gaussian $N(A^{-1}b, A^{-1})$ is equivalent to solving the linear system $Ax = b$ (see [17, 18, 19]), our notation aligns with literature on solving linear systems.

Chapter 3

Gaussian targets

3.1 Expected acceptance rate for a Gaussian target

The expected acceptance rate is a quantity that is related to efficiency and for optimal performance the proposal is usually tuned so that the observed average acceptance rate is between 0 and 1. For example, it has been shown, for particular target distributions, in the case when $d \rightarrow \infty$, that 0.234 is optimal for RWM [29], 0.574 is optimal for MALA [30] and SLA [9], and 0.651 is optimal for HMC [6]. All of these results required expressions for the expected acceptance rate of the algorithm as $d \rightarrow \infty$. Here we derive an expression for the expected acceptance rate for a MH algorithm with an AR(1) proposal (1.5) and Gaussian target $N(A^{-1}b, A^{-1})$, provided the splitting matrices are functions of A . Thus, our MH algorithm is defined by

$$\begin{aligned} \text{Target:} & \quad N(A^{-1}b, A^{-1}), \\ \text{Proposal:} & \quad y = Gx + M^{-1}\beta + (\mathcal{A}^{-1} - G\mathcal{A}^{-1}G^T)^{1/2}\xi, \end{aligned} \tag{3.1}$$

where $\xi \sim N(0, I)$, and we have used $G = M^{-1}N$ and $M^{-1}(M^T + N)M^{-T} = \mathcal{A}^{-1} - G\mathcal{A}^{-1}G^T$ [18, Lem. 2.3]. The following lemma is a result of simple algebra. The proof is in the Appendix.

Lemma 3.1.1 Suppose $\mathcal{A} = M - N$ is a symmetric splitting. Then the acceptance probability for (3.1) satisfies

$$\alpha(x, y) = 1 \wedge \exp\left(-\frac{1}{2}y^T(A - \mathcal{A})y + \frac{1}{2}x^T(A - \mathcal{A})x + (b - \beta)^T(y - x)\right).$$

Since A is real and symmetric, we can define a spectral decomposition

$$A = Q\Lambda Q^T$$

where $Q \in \mathbb{R}^{d \times d}$ is orthogonal and $\Lambda = \text{diag}(\lambda_1^2, \dots, \lambda_d^2)$ is a diagonal matrix of eigenvalues of A . Although we may not be able to compute Q and Λ this does not stop us from using the existence of a spectral decomposition for theory. Simple algebra gives us the following result.

Lemma 3.1.2 Suppose $M = M(A)$ and $N = N(A)$ are functions of A . Then G and \mathcal{A} are also functions of A and under the coordinate transformation

$$x \leftrightarrow Q^T x$$

the MH algorithm (3.1) is transformed to a MH algorithm defined by

$$\begin{aligned} \text{Target:} & \quad N(\Lambda^{-1}Q^T b, \Lambda^{-1}), \\ \text{Proposal:} & \quad y = Gx + M(\Lambda)^{-1}Q^T\beta + (\tilde{\Lambda}^{-1} - G\tilde{\Lambda}^{-1}G^T)^{1/2}\xi, \end{aligned} \tag{3.2}$$

where $\xi \sim N(0, I)$, and $G = M(\Lambda)^{-1}N(\Lambda)$ and $\tilde{\Lambda} = \mathcal{A}(\Lambda)$ are diagonal matrices.

Using Lemma 3.1.1 we see that the acceptance probability of MH algorithms (3.1) and (3.2) are identical and hence it is sufficient to analyse the convergence properties of (3.2) to determine the convergence properties of (3.1).

We will use the following Lyapunov central limit theorem, see e.g. [12, Thm. 27.3].

Theorem 3.1.1 For each $d \in \mathbb{N}$ let $X_{d,1}, \dots, X_{d,d}$ be a sequence of independent random variables each with finite expected value $\mu_{d,i}$ and variance $\sigma_{d,i}^2$. Define $s_d^2 := \sum_{i=1}^d \sigma_{d,i}^2$. If there exists a $\delta > 0$ such that

$$\lim_{d \rightarrow \infty} \frac{1}{s_d^{2+\delta}} \sum_{i=1}^d \mathbb{E}[|X_{d,i} - \mu_{d,i}|^{2+\delta}] = 0,$$

then

$$\frac{1}{s_d} \sum_{i=1}^d (X_{d,i} - \mu_{d,i}) \xrightarrow{\mathcal{D}} \mathcal{N}(0, 1) \quad \text{as } d \rightarrow \infty.$$

An equivalent conclusion to this theorem is $\sum_{i=1}^d X_{d,i} \rightarrow \mathcal{N}(\sum_{i=1}^d \mu_{d,i}, s_d^2)$ in distribution as $d \rightarrow \infty$. Another useful fact is

$$X \sim \mathcal{N}(\mu, \sigma^2) \quad \Rightarrow \quad \mathbb{E}[1 \wedge e^X] = \Phi\left(\frac{\mu}{\sigma}\right) + e^{\mu+\sigma^2/2} \Phi\left(-\sigma - \frac{\mu}{\sigma}\right) \quad (3.3)$$

where Φ is the standard normal cumulative distribution function. See e.g. [29, Prop. 2.4] or [9, Lem. B.2].

Theorem 3.1.2 Suppose that M and N in (3.1) are functions of A , and the Markov chain is in equilibrium, i.e. $x \sim \mathcal{N}(A^{-1}b, A^{-1})$. If there exists a $\delta > 0$ such that

$$\lim_{d \rightarrow \infty} \frac{\sum_{i=1}^d |T_{ji}|^{2+\delta}}{\left(\sum_{i=1}^d |T_{ji}|^2\right)^{1+\delta/2}} = 0 \quad \text{for } j = 1, 2, 3, 4, 5 \quad (3.4)$$

($j = 0$ is not required) and the limits $\mu = \lim_{d \rightarrow \infty} \sum_{i=1}^d \mu_{d,i}$ and $\sigma^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d \sigma_{d,i}^2$ exist where

$$\mu_{d,i} = T_{0i} + T_{3i} + T_{4i} \quad \text{and} \quad \sigma_{d,i}^2 = T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2,$$

then

$$Z := \log \left(\frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} \right) \xrightarrow{\mathcal{D}} \mathcal{N}(\mu, \sigma^2) \quad \text{as } d \rightarrow \infty$$

and the expected acceptance probability satisfies

$$\mathbb{E}[\alpha(x, y)] = \mathbb{E}[1 \wedge e^Z] \rightarrow \Phi\left(\frac{\mu}{\sigma}\right) + e^{\mu+\sigma^2/2} \Phi\left(-\sigma - \frac{\mu}{\sigma}\right) \quad \text{as } d \rightarrow \infty.$$

In the above theorem, M and N may depend on d , so T_{ji} may depend on d .

Proof. By Lemma 3.1.2 it is sufficient to only consider (3.1) in the case where all matrices are diagonal matrices, e.g. $A = \text{diag}(\lambda_1^2, \dots, \lambda_d^2)$, $\mathcal{A} = \text{diag}(\tilde{\lambda}_1^2, \dots, \tilde{\lambda}_d^2)$, $M = \text{diag}(M_1, \dots, M_d)$, $G = \text{diag}(G_1, \dots, G_d)$, $m_i = \lambda_i^{-2} b_i$, and $\tilde{m}_i = \tilde{\lambda}_i^{-2} \beta_i$. Then, in equilibrium we have

$$x_i = m_i + \frac{1}{\lambda_i} \xi_i$$

where $\xi_i \sim \mathcal{N}(0, 1)$ and using $\tilde{m}_i = G_i \tilde{m}_i + M_i^{-1} \beta_i$ we have

$$\begin{aligned} y_i &= G_i x_i + \frac{\beta_i}{M_i} + \frac{(1 - G_i^2)^{1/2}}{\tilde{\lambda}_i} \nu_i \\ &= G_i \left(m_i + \frac{1}{\lambda_i} \xi_i \right) + (1 - G_i) \tilde{m}_i + \frac{g_i^{1/2}}{\tilde{\lambda}_i} \nu_i \\ &= \tilde{m}_i + G_i \hat{r} + \frac{G_i}{\lambda_i} \xi_i + \frac{g_i^{1/2}}{\tilde{\lambda}_i} \nu_i \end{aligned}$$

where $\nu_i \sim \mathcal{N}(0, 1)$. From Lemma 3.1.1 we also have $Z = \sum_{i=1}^d Z_{d,i}$ where

$$Z_{d,i} = -\frac{1}{2}(\lambda_i^2 - \tilde{\lambda}_i^2)(y_i^2 - x_i^2) + (b_i - \beta_i)(y_i - x_i).$$

Substituting x_i and y_i as above, using the identity $(b_i - \beta_i)\lambda_i^{-2} = \hat{r}_i + r_i \tilde{m}_i$, then after some algebra we eventually find

$$Z_{d,i} = T_{0i} + T_{1i}\xi_i + T_{2i}\nu_i + T_{3i}\xi_i^2 + T_{4i}\nu_i^2 + T_{5i}\xi_i\nu_i.$$

Hence

$$\mu_{d,i} := \mathbb{E}[Z_{d,i}] = T_{0i} + T_{3i} + T_{4i}$$

and

$$\begin{aligned} \sigma_{d,i}^2 &:= \text{Var}[Z_{d,i}] = \mathbb{E}[Z_{d,i}^2] - \mathbb{E}[Z_{d,i}]^2 \\ &= (T_{0i}^2 + T_{1i}^2 + T_{2i}^2 + 3T_{3i}^2 + 3T_{4i}^2 + T_{5i}^2 + 2T_{0i}T_{3i} + 2T_{0i}T_{4i} + 2T_{3i}T_{4i}) \\ &\quad - (T_{0i} + T_{3i} + T_{4i})^2 \\ &= T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 \end{aligned}$$

and

$$Z_{d,i} - \mu_{d,i} = T_{1i}\xi_i + T_{2i}\nu_i + T_{3i}(\xi_i^2 - 1) + T_{4i}(\nu_i^2 - 1) + T_{5i}\xi_i\nu_i.$$

Therefore, for any $d \in \mathbb{N}$ and $\delta > 0$ we can bound the Lyapunov condition in Theorem 3.1.1 as follows

$$\begin{aligned} \frac{1}{s_d^{2+\delta}} \sum_{i=1}^d \mathbb{E}[|Z_{d,i} - \mu_{d,i}|^{2+\delta}] &\leq \frac{5^{2+\delta}}{s_d^{2+\delta}} \sum_{j=1}^5 C_j(\delta) \sum_{i=1}^d |T_{ji}|^{2+\delta} \\ &\leq 5^{2+\delta} \sum_{j=1}^5 C_j(\delta) \frac{\sum_{i=1}^d |T_{ji}|^{2+\delta}}{\left(\sum_{i=1}^d T_{ji}^2\right)^{1+\delta/2}} \end{aligned}$$

where $C_1(\delta) = C_2(\delta) = \mathbb{E}[|\xi|^{2+\delta}]$, $C_3(\delta) = C_4(\delta) = \mathbb{E}[|\xi^2 - 1|^{2+\delta}]$ and $C_5(\delta) = \mathbb{E}[|\xi|^{2+\delta}]^2$, and $\xi \sim \mathcal{N}(0, 1)$.

Therefore, if (3.4) holds then the result follows from Theorem 3.1.1 and (3.3). \blacksquare

3.2 Expected squared jump size for a Gaussian target

The efficiency of a MCMC method is usually given in terms of the integrated autocorrelation time which is equivalent to “the number of dependent sample points from the Markov chain needed to give the variance reducing power of one independent point”, see e.g. [26, §6.3]. Unfortunately, we are unable to directly estimate this quantity for our matrix splitting methods, and it depends on the statistic of concern. As a proxy we instead consider the expected squared jump size of the Markov chain in a direction $q \in \mathbb{R}^d$,

$$\mathbb{E}[(q^T(x' - x))^2]$$

where $x, x' \sim N(A^{-1}b, A^{-1})$ are successive elements of the Markov chain in equilibrium. We will only consider the cases where q is an eigenvector of the precision or covariance matrix. It is related to the integrated autocorrelation time for the linear functional $q^T(\cdot)$ by

$$\text{Corr}[q^T x, q^T x'] = 1 - \frac{\mathbb{E}[(q^T(x' - x))^2]}{2\text{Var}[q^T x]}$$

so that large squared jump size implies small first-order autocorrelation, see e.g. [30, §3] or [9, §2.3].

This is similar to the approach used for analysing the efficiency of RWM, MALA and HMC, where the expected squared jump size of an arbitrary component of the Markov chain is considered, see e.g. [9] and [6].

We will need the following technical lemma whose proof is in the Appendix.

Lemma 3.2.1 Suppose $\{t_i\}_{i=1}^\infty \subset \mathbb{R}$ and $r > 0$. Then, for any $k \in \mathbb{N}$,

$$\lim_{d \rightarrow \infty} \frac{\sum_{i=1}^d |t_i|^r}{\left(\sum_{i=1}^d t_i^2\right)^{r/2}} = 0 \quad \Rightarrow \quad \lim_{d \rightarrow \infty} \frac{\sum_{i=1, i \neq k}^d |t_i|^r}{\left(\sum_{i=1, i \neq k}^d t_i^2\right)^{r/2}} = 0. \quad (3.5)$$

The following theorem is a generalization of [6, Prop. 3.8].

Theorem 3.2.1 Suppose that M and N in (3.1) are functions of A , and the Markov chain is in equilibrium, i.e. $x \sim N(A^{-1}b, A^{-1})$. With $\mu_{d,i}$ and $\sigma_{d,i}^2$ defined as in Theorem 3.1.2, let q_i be a normalized eigenvector of A corresponding to λ_i^2 . If there exists a $\delta > 0$ such that (3.4) is satisfied, and $\mu^- := \lim_{d \rightarrow \infty} \sum_{j=1, j \neq i}^d \mu_{d,j}$ and $(\sigma^-)^2 := \lim_{d \rightarrow \infty} \sum_{j=1, j \neq i}^d \sigma_{d,j}^2$ exist, then

$$\mathbb{E}[(q_i^T(x' - x))^2] = U_1 U_2 + E_3 + o(U_1) \quad (3.6)$$

as $d \rightarrow \infty$ where

$$\begin{aligned} U_1 &= \tilde{g}_i^2 \hat{r}_i^2 + \frac{\tilde{g}_i^2}{\lambda_i^2} + \frac{g_i}{\tilde{\lambda}_i^2}, \\ U_2 &= \mathbb{E}[1 \wedge e^X] = \Phi\left(\frac{\mu^-}{\sigma^-}\right) + e^{\mu^- + (\sigma^-)^2/2} \Phi\left(-\sigma^- - \frac{\mu^-}{\sigma^-}\right), \quad X \sim N(\mu^-, (\sigma^-)^2), \\ |E_3| \leq U_3 &= (\sigma_{d,i}^2 + \mu_{d,i}^2)^{1/2} \times \left(\tilde{g}_i^4 \hat{r}_i^4 + \frac{3}{\lambda_i^4} (\tilde{g}_i^2 + \tilde{r}_i g_i)^2 + \frac{6}{\lambda_i^2} \hat{r}_i^2 \tilde{g}_i^2 (\tilde{g}_i^2 + \tilde{r}_i g_i) \right)^{1/2}. \end{aligned}$$

Proof. Under the coordinate transformation $x \leftrightarrow Q^T x$, (3.1) becomes (3.2) and $\mathbb{E}[(q_i^T(x' - x))^2]$ becomes $\mathbb{E}[(x'_i - x_i)^2]$. Therefore it is sufficient to only consider the squared jump size of an arbitrary coordinate of the Markov chain for the case when all matrices are diagonal matrices. As in the proof of Theorem 3.1.2, let $A = \text{diag}(\lambda_1^2, \dots, \lambda_d^2)$, $\mathcal{A} = \text{diag}(\tilde{\lambda}_1^2, \dots, \tilde{\lambda}_d^2)$, $M = \text{diag}(M_1, \dots, M_d)$, $G = \text{diag}(G_1, \dots, G_d)$, $m_i = \lambda_i^{-2} b_i$, and $\tilde{m}_i = \tilde{\lambda}_i^{-2} \beta_i$. Since the chain is in equilibrium we have $x_i = m_i + \lambda_i^{-1} \xi_i$ for $\xi_i \sim \mathcal{N}(0, 1)$ and $y_i = \tilde{m}_i + G_i \hat{r} + G_i \lambda_i^{-1} \xi_i + g_i^{1/2} \tilde{\lambda}_i^{-1} \nu_i$ where $\nu_i \sim \mathcal{N}(0, 1)$. Define $\alpha^-(x, y) := 1 \wedge \exp(\sum_{j=1, j \neq i}^d Z_{d,j})$ where $Z_{d,j}$ is defined as in the proof of Theorem 3.1.2.

The proof strategy is now to approximate $\mathbb{E}[(x'_i - x_i)^2] = \mathbb{E}[(y_i - x_i)^2 \alpha(x, y)]$ with $\mathbb{E}[(y_i - x_i)^2 \alpha^-(x, y)]$;

$$\mathbb{E}[(x'_i - x_i)^2] = \mathbb{E}[(y_i - x_i)^2 \alpha^-(x, y)] + \mathbb{E}[(\alpha(x, y) - \alpha^-(x, y))(y_i - x_i)^2].$$

By independence,

$$\begin{aligned} \mathbb{E}[(y_i - x_i)^2 \alpha^-(x, y)] &= \mathbb{E}[(y_i - x_i)^2] \mathbb{E}[\alpha^-(x, y)] \\ &= \mathbb{E} \left[\left(-\tilde{g}_i \hat{r}_i - \frac{\tilde{g}_i}{\tilde{\lambda}_i} \xi_i + \frac{g^{1/2}}{\tilde{\lambda}_i} \nu_i \right)^2 \right] \mathbb{E}[\alpha^-(x, y)] \\ &= U_1 \mathbb{E}[\alpha^-(x, y)]. \end{aligned}$$

Also, by Theorem 3.1.2 (using Lemma 3.2.1 to ensure the appropriate condition for Theorem 3.1.2 is met) we obtain $\mathbb{E}[\alpha^-(x, y)] \rightarrow U_2$ as $d \rightarrow \infty$, so $\mathbb{E}[(y_i - x_i)^2 \alpha(x, y)] = U_1 U_2 + o(U_1)$ as $d \rightarrow \infty$.

The error is bounded using the Cauchy-Schwarz inequality;

$$|\mathbb{E}[(\alpha(x, y) - \alpha^-(x, y))(y_i - x_i)^2]| \leq \mathbb{E}[(\alpha(x, y) - \alpha^-(x, y))^2]^{1/2} \mathbb{E}[(y_i - x_i)^4]^{1/2}.$$

Since $1 \wedge e^X$ is Lipschitz with constant 1, and using results from the proof of Theorem 3.1.2, we obtain

$$\mathbb{E}[(\alpha(x, y) - \alpha^-(x, y))^2]^{1/2} \leq \mathbb{E}[Z_{d,i}^2]^{1/2} = (\sigma_{d,i}^2 + \mu_{d,i}^2)^{1/2},$$

and some algebra yields

$$\begin{aligned} \mathbb{E}[(y_i - x_i)^4]^{1/2} &= \mathbb{E} \left[\left(-\tilde{g}_i \hat{r}_i - \frac{\tilde{g}_i}{\tilde{\lambda}_i} \xi_i + \frac{g^{1/2}}{\tilde{\lambda}_i} \nu_i \right)^4 \right]^{1/2} \\ &= \left(\tilde{g}_i^4 \hat{r}_i^4 + \frac{3}{\tilde{\lambda}_i^4} (\tilde{g}_i^2 + \tilde{r}_i g_i)^2 + \frac{6}{\tilde{\lambda}_i^2} \hat{r}_i^2 \tilde{g}_i^2 (\tilde{g}_i^2 + \tilde{r}_i g_i) \right)^{1/2}. \end{aligned}$$

■

The terms in the theorem above are quite lengthy, but in many situations they simplify. For example, we may have the situation where $\mu^- \rightarrow \mu$ and $\sigma^- \rightarrow \sigma$ as $d \rightarrow \infty$, so U_2 becomes the expected acceptance rate for the algorithm. Also, it may be possible to derive a bound such as

$$|U_3| \leq C(r_i + \hat{r}_i)$$

so that U_3 is small if both the relative error of the i^{th} eigenvalue and error of the means are small.

Chapter 4

Non-Gaussian targets

Our results can be extended to non-Gaussian target distributions in some cases. We follow the methodology in [9]. Suppose that target π_d is a change of measure from a reference Gaussian, so that π_d is defined by (1.2) and (1.4).

With an AR(1) proposal associated with a matrix splitting $\mathcal{A} = M - N$, we consider an MH algorithm defined by

$$\begin{aligned} \text{Target:} & \quad \pi_d, \\ \text{Proposal:} & \quad y = Gx + g + \nu, \quad \text{where } \nu \sim \text{N}(0, \Sigma), \end{aligned} \tag{4.1}$$

where $G = M^{-1}N$, $g = M^{-1}\beta$ and $\Sigma = M^{-1}(M^T + N)M^{-T} = \mathcal{A}^{-1} - G\mathcal{A}^{-1}G^T$. The acceptance probability of this MH algorithm satisfies

$$\alpha(x, y) = 1 \wedge \exp(\phi_d(x) - \phi_d(y) + Z)$$

where $Z = \log\left(\frac{\tilde{\pi}_d(y)q(y,x)}{\tilde{\pi}_d(x)q(x,y)}\right)$. Define $\tilde{\alpha}(x, y) = 1 \wedge \exp(Z)$ to be the acceptance probability for MH algorithm (3.1).

In the theory below, $\mathbb{E}_{\pi_d}[\alpha(x, y)]$ is the expectation of $\alpha(x, y)$ over $x \sim \pi_d$ and y from (4.1).

4.1 Expected acceptance rate for a non-Gaussian target

The following theorem applies to inverse problems with a Gaussian prior and bounded likelihood.

Theorem 4.1.1 Suppose there exists a constant $M > 0$ such that for sufficiently large d

$$|\phi_d(x)| \leq M \quad \text{for all } x \in \mathbb{R}^d.$$

Then MH algorithm (4.1) in equilibrium satisfies

$$\mathbb{E}_{\pi_d}[\alpha(x, y)] \leq C\mathbb{E}_{\tilde{\pi}_d}[\tilde{\alpha}(x, y)] \tag{4.2}$$

$$\mathbb{E}_{\pi_d}[\alpha(x, y)] > 0 \quad \text{if } \mathbb{E}_{\tilde{\pi}_d}[|Z|] < \infty, \tag{4.3}$$

where $\mathbb{E}_{\tilde{\pi}_d}[\tilde{\alpha}(x, y)]$ is the expected acceptance rate of (3.1) in equilibrium.

Proof. We follow the same reasoning as in the proof of [9, Thm. 2]. Note that $1 \wedge \exp(\phi_d(x) - \phi_d(y) + Z) \leq \exp(2M)(1 \wedge \exp(Z))$, and $\pi_d(x) \leq \exp(M)\tilde{\pi}_d(x)$. Hence, we obtain (4.2) with $C = \exp(3M)$.

To prove (4.3) first note for a random variable X and any $\gamma > 0$ we have $\mathbb{E}[1 \wedge \exp(X)] \geq \exp(-\gamma)(1 - \gamma^{-1}\mathbb{E}[|X|])$, see [9, Lem. B1]. Also note that $C_0 := \mathbb{E}_{\pi_d}[|\phi_d(x) - \phi_d(y) + Z|] \leq C + C\mathbb{E}_{\tilde{\pi}_d}[|Z|] < \infty$. Hence, we obtain (4.3) by taking $\gamma = 2C_0$. \blacksquare

Thus, in a certain weak sense, the acceptance rate of (4.1) with non-Gaussian target mimics the acceptance rate of (3.1) with a Gaussian target; if the acceptance rate of (3.1) is small, then so is the acceptance rate of (4.1); and if the expected value of $|Z|$ is finite (which loosely corresponds to when the acceptance rate of (3.1) is positive) then the acceptance rate of (4.1) in equilibrium is positive.

A similar result is given in [9, Thm. 2] for RWM and SLA.

Our next theorem more precisely describes the acceptance rate for a non-Gaussian target, but first, some definitions and a lemma.

We associate a norm with the spd precision matrix A of our reference Gaussian measure $\tilde{\pi}$. For any $s \in \mathbb{R}$ define a norm $|\cdot|_s$ on \mathbb{R}^d by

$$|x|_s = |A^s x|$$

for all $x \in \mathbb{R}^d$. If λ_1^2 is the smallest eigenvalue of A , then

$$|x|_s \leq \lambda_1^{2(s-r)} |x|_r \quad \text{for all } s < r. \quad (4.4)$$

Assumption 4.1.2 Suppose there exist constants $m, s, s', s'' \in \mathbb{R}$, $C, p > 0$, and a locally bounded function $\delta : \mathbb{R}^+ \times \mathbb{R}^+ \mapsto \mathbb{R}^+$ such that for all sufficiently large d

$$\begin{aligned} \phi_d(x) &\geq m, \\ |\phi_d(x) - \phi_d(y)| &\leq \delta(|x - A^{-1}b|_s, |y - A^{-1}b|_s) |x - y|_{s'}, \\ |\phi_d(x)| &\leq C(1 + |x - A^{-1}b|_{s''}^p) \end{aligned}$$

for all $x, y \in \mathbb{R}^d$.

Assumption 4.1.3 Suppose that r is such that

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d \lambda_i^{4r-1} < \infty.$$

The following lemma is similar to part of the proof of [9, Thm. 3].

Lemma 4.1.1 Suppose ϕ_d satisfies Assumption 4.1.2 and $r = \max\{s, s', s''\}$ satisfies Assumption 4.1.3. Also suppose that there exists $t_{d,i}$ and a $t > 0$ such that MH algorithm (4.1) satisfies

$$\begin{aligned} G \text{ and } \Sigma &\text{ are functions of } A, \\ \tilde{g}_i^2 \tilde{r}_i^2 \lambda_i, \tilde{g}_i^2 \lambda_i^{-1}, g_i \tilde{\lambda}_i^{-1} &\text{ are } \mathcal{O}(t_{d,i}) = \mathcal{O}(d^{-t}) \text{ (uniformly in } i) \text{ as } d \rightarrow \infty, \text{ and} \\ \tilde{r}_i &\text{ is bounded uniformly in } d \text{ and } i. \end{aligned}$$

Then for any $q \in \mathbb{N}$ there exists a constant $C > 0$ (that may depend on q) such that

$$\begin{aligned} \mathbb{E}_{\tilde{\pi}_d}[|x - A^{-1}b|_r^{2q}] &< C \text{ for all } d \\ \mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^{2q}] &= \mathcal{O}(t_{d,i}^q) = \mathcal{O}(d^{-qt}) \text{ as } d \rightarrow \infty, \end{aligned}$$

and for proposal y from x ,

$$\phi_d(x) - \phi_d(y) \rightarrow 0 \quad \text{in } L^q(\pi_d) \text{ as } d \rightarrow \infty.$$

Proof. For $x \sim N(A^{-1}b, A^{-1})$, $\xi \sim N(0, I)$, $\nu \sim N(0, I)$ and $\Lambda = \text{diag}(\lambda_1^2, \dots, \lambda_d^2)$,

$$\begin{aligned} \mathbb{E}_{\tilde{\pi}_d}[|x - A^{-1}b|_r^{2q}] &= \mathbb{E}_{\tilde{\pi}_d}[|A^{r-1/2}\xi|^{2q}] = \mathbb{E}_{\tilde{\pi}_d}[|\Lambda^{r-1/2}\nu|^{2q}] \\ &= \mathbb{E}_{\tilde{\pi}_d} \left[\left(\sum_{i=1}^d \lambda_i^{4r-1} \nu_i^2 \right)^q \right] \leq C \left(\sum_{i=1}^d \lambda_i^{4r-1} \right)^q, \end{aligned}$$

which is bounded uniformly in d by Assumption 4.1.3.

As above, and using the transformation $x \leftrightarrow Q^T x$, and $(y_i - x_i) = -\tilde{g}_i \hat{r}_i - \tilde{g}_i \lambda_i^{-1} \xi_i + g_i^{1/2} \tilde{\lambda}_i^{-1} \nu_i$ from the proof of Theorem 3.2.1 where ξ_i and $\nu_i \stackrel{\text{iid}}{\sim} N(0, 1)$, we have

$$\begin{aligned} \mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^{2q}] &= \mathbb{E}_{\tilde{\pi}_d} \left[\left(\sum_{i=1}^d \lambda_i^{4r} \left(-\tilde{g}_i \hat{r}_i - \frac{\tilde{g}_i}{\lambda_i} \xi_i + \frac{g_i^{1/2}}{\tilde{\lambda}_i} \nu_i \right)^2 \right)^q \right] \\ &= \mathbb{E}_{\tilde{\pi}_d} \left[\left(\sum_{i=1}^d \lambda_i^{4r-1} \left(-\tilde{g}_i \hat{r}_i \lambda_i^{1/2} - \frac{\tilde{g}_i}{\lambda_i^{1/2}} \xi_i + \frac{\tilde{r}_i^{1/4} g_i^{1/2}}{\tilde{\lambda}_i^{1/2}} \nu_i \right)^2 \right)^q \right]. \end{aligned}$$

Since $\tilde{g}_i^2 \hat{r}_i^2 \lambda_i$, $\tilde{g}_i^2 \lambda_i^{-1}$, $g_i \tilde{\lambda}_i^{-1} = \mathcal{O}(t_{d,i})$ uniformly in i , and \tilde{r}_i is bounded uniformly in d and i , it follows that for all sufficiently large d ,

$$\mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^{2q}] \leq C t_{d,i}^q \left(\sum_{i=1}^d \lambda_i^{2r-1} \right)^q,$$

so by Assumption 4.1.3, $\mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^{2q}] = \mathcal{O}(t_{d,i}^q)$ as $d \rightarrow \infty$.

From $\mathbb{E}_{\tilde{\pi}_d}[|x - A^{-1}b|_r^{2q}] < C$ and $\mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^{2q}] \rightarrow 0$, it follows from the triangle inequality that there is a (new) constant $C > 0$ such that

$$\mathbb{E}_{\tilde{\pi}_d}[|y - A^{-1}b|_r^{2q}] < C \quad \text{for all } d. \quad (4.5)$$

Let $\Delta_d = \phi_d(x) - \phi_d(y)$. For any $R > 0$ define

$$\begin{aligned} \gamma(R) &= \sup\{\delta(a, b)^q : a \leq R, b \leq R\} \\ S_1 &= \{x \in \mathbb{R}^d : |x - A^{-1}b|_s \leq R\}, \\ S_2 &= \{x \in \mathbb{R}^d : |y - A^{-1}b|_s \leq R\}, \end{aligned}$$

and let \mathbb{I}_S be the indicator function for set S . Using Assumption 4.1.2, a generic constant C that may vary between lines, the Cauchy-Schwarz inequality, and then Markov's inequality, we have for each $q \in \mathbb{N}$

$$\begin{aligned} \mathbb{E}_{\tilde{\pi}_d}[|\Delta_d|^q] &= \mathbb{E}_{\tilde{\pi}_d}[|\Delta_d|^q \mathbb{I}_{S_1 \cap S_2}] + \mathbb{E}_{\tilde{\pi}_d}[|\Delta_d|^q \mathbb{I}_{\mathbb{R}^d \setminus (S_1 \cap S_2)}] \\ &\leq \gamma(R) \mathbb{E}[|x - y|_{s'}^q] + C \mathbb{E}[(1 + |x - A^{-1}b|_{s''}^{pq} + |y - A^{-1}b|_{s''}^{pq}) \mathbb{I}_{\mathbb{R}^d \setminus (S_1 \cap S_2)}] \\ &\leq \gamma(R) \mathbb{E}[|x - y|_{s'}^q] \\ &\quad + C \mathbb{E}[1 + |x - A^{-1}b|_{s''}^{2pq} + |y - A^{-1}b|_{s''}^{2pq}]^{1/2} (\mathbb{P}(\mathbb{R}^d \setminus S_1) + \mathbb{P}(\mathbb{R}^d \setminus S_2))^{1/2} \\ &\leq \gamma(R) \mathbb{E}[|x - y|_{s'}^q] + C (\mathbb{P}(\mathbb{R}^d \setminus S_1) + \mathbb{P}(\mathbb{R}^d \setminus S_2))^{1/2} \\ &\leq \gamma(R) \mathbb{E}[|x - y|_{s'}^q] + \frac{C}{R^{1/2}} (\mathbb{E}[|x - A^{-1}b|_s] + \mathbb{E}[|y - A^{-1}b|_s])^{1/2} \\ &\leq \gamma(R) \mathbb{E}[|x - y|_{s'}^q] + \frac{C}{R^{1/2}}. \end{aligned}$$

Note that we used Jensen's inequality (which implies $\|f\|_{L^p(\tilde{\pi}_d)} \leq \|f\|_{L^q(\tilde{\pi}_d)}$ for $1 \leq p \leq q \leq \infty$) and (4.4) to obtain bounds on $\mathbb{E}[|x - A^{-1}b|_{s''}^{2pq}]$, $\mathbb{E}[|y - A^{-1}b|_{s''}^{2pq}]$, $\mathbb{E}[|x - A^{-1}b|_s]$ and $\mathbb{E}[|y - A^{-1}b|_s]$.

Hence, for any $\epsilon > 0$ we can choose $R = R(\epsilon)$ such that $C/R^{1/2} < \epsilon/2$ and since $\mathbb{E}_{\tilde{\pi}_d}[|y - x|_r^q] \rightarrow 0$ as $d \rightarrow \infty$ (by Jensen's inequality), for all sufficiently large d we have

$$\mathbb{E}_{\tilde{\pi}_d}[|\phi_d(x) - \phi_d(y)|] = \mathbb{E}_{\tilde{\pi}_d}[|\Delta_d|^q] < \epsilon.$$

Thus,

$$\phi_d(x) - \phi_d(y) \rightarrow 0 \quad \text{in } L^q(\tilde{\pi}_d) \text{ as } d \rightarrow \infty.$$

The result then follows from $\phi_d(x) \geq m$. \blacksquare

Theorem 4.1.4 Suppose that ϕ_d satisfies Assumption 4.1.2 and $r = \max\{s, s', s''\}$ satisfies Assumption 4.1.3. Also suppose there exists $t_{d,i}$ and $t > 0$ such that MH algorithm (4.1) satisfies

G and Σ are functions of A ,

$\tilde{g}_i^2 \hat{r}_i^2 \lambda_i$, $\tilde{g}_i^2 \lambda_i^{-1}$, $g_i \tilde{\lambda}_i^{-1}$ are $\mathcal{O}(t_{d,i}) = \mathcal{O}(d^{-t})$ (uniformly in i) as $d \rightarrow \infty$,

\tilde{r}_i is bounded uniformly in d and i ,

$T_{1i}, T_{2i}, T_{3i}, T_{4i}, T_{5i}$ are $\mathcal{O}(d^{-1/2})$ as $d \rightarrow \infty$ (uniformly in i), and

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 < \infty.$$

If

$$\begin{aligned} \mu_{ng} &= \mu + \lim_{d \rightarrow \infty} \sum_{i=1}^d \kappa_{d,i} T_{1i} + T_{3i}(\gamma_i - 1), \\ \sigma_{ng}^2 &= \sigma^2 + \lim_{d \rightarrow \infty} \sum_{i=1}^d (\kappa_{d,i} T_{1i} + T_{3i}(\gamma_{d,i} - 1))^2, \end{aligned}$$

exist, where $\kappa_{d,i} = \mathbb{E}_{\pi_d}[q_i^T A^{1/2}(x - A^{-1}b)]$, $\gamma_{d,i} = \mathbb{E}_{\pi_d}[(q_i^T A^{1/2}(x - A^{-1}b))^2]$, q_i is a normalised eigenvector of A corresponding to the eigenvalue λ_i^2 , and μ and σ^2 are the same as in Theorem 3.1.2, then

$$\mathbb{E}_{\pi_d}[\alpha(x, y)] \rightarrow \mathbb{E}[1 \wedge e^{Z_{ng}}] = \Phi\left(\frac{\mu_{ng}}{\sigma_{ng}}\right) + e^{\mu_{ng} + \sigma_{ng}^2/2} \Phi\left(-\sigma_{ng} - \frac{\mu_{ng}}{\sigma_{ng}}\right)$$

as $d \rightarrow \infty$, where $Z_{ng} \sim \mathcal{N}(\mu_{ng}, \sigma_{ng}^2)$.

Proof. As in the proofs of Theorems 3.1.2 and 3.2.1 it is sufficient to prove the result in the case when all matrices are diagonal. This follows from the coordinate transformation $z = Q^T x$ where $A = Q\Lambda Q^T$, since

$$|x|_s = |\Lambda^s z| \quad \text{and} \quad |x - A^{-1}b|_s = |\Lambda^s(z - \Lambda^{-1}Q^T b)|$$

for all $x \in \mathbb{R}^d$ and $s \in \mathbb{R}$, and since $\psi_d(z) := \phi_d(Qz)$ satisfies Assumption 4.1.2. Henceforth and without loss of generality, let us assume that A , G and Σ are diagonal matrices.

Using Lemma 4.1.1 with $q = 1$, and the fact that $z \mapsto 1 \wedge \exp(z)$ is globally Lipschitz continuous, it follows that

$$\mathbb{E}_{\pi_d}[\alpha(x, y)] - \mathbb{E}_{\pi_d}[\tilde{\alpha}(x, y)] \rightarrow 0 \quad \text{as } d \rightarrow \infty. \quad (4.6)$$

To complete the proof we must find the limit of $\mathbb{E}_{\pi_d}[\tilde{\alpha}(x, y)]$ as $d \rightarrow \infty$.

As in the proof of Theorem 3.1.2 we have $Z = \sum_{i=1}^d Z_{d,i}$ where

$$Z_{d,i} = T_{0i} + T_{1i}\xi_i + T_{2i}\nu_i + T_{3i}\xi_i^2 + T_{4i}\nu_i^2 + T_{5i}\xi_i\nu_i,$$

noting that $\xi_i = \lambda_i(x_i - m_i)$ with $x \sim \pi_d$, $m_i = (A^{-1}b)_i$, and $\nu_i \sim \mathcal{N}(0, 1)$.

Note that $\kappa_{d,i} = \mathbb{E}_{\pi_d}[\lambda_i(x_i - m_i)] = \mathbb{E}_{\pi_d}[\xi_i]$ and $\gamma_{d,i} = \mathbb{E}_{\pi_d}[\lambda_i^2(x_i - m_i)^2] = \mathbb{E}_{\pi_d}[\xi_i^2]$ are both uniformly bounded in d and i since $\phi_d(x) \geq m$ and $|\kappa_{d,i}| \leq e^{-m}\mathbb{E}_{\tilde{\pi}_d}[|\xi_i|] = e^{-m}\mathbb{E}[|u|]$ and $0 \leq \gamma_{d,i} \leq e^{-m}\mathbb{E}_{\tilde{\pi}_d}[\xi_i^2] = e^{-m}\mathbb{E}[u^2]$ where $u \sim \mathcal{N}(0, 1)$. If we define

$$S_{d,j} := \sum_{i=1}^j T_{1i}(\xi_i - \kappa_{d,i}) + T_{2i}\nu_i + T_{3i}(\xi_i^2 - \gamma_{d,i}) + T_{4i}(\nu_i^2 - 1) + T_{5i}\xi_i\nu_i,$$

then

$$Z = \sum_{i=1}^d (T_{0i} + T_{1i}\kappa_{d,i} + T_{3i}\gamma_{d,i} + T_{4i}) + S_{d,d}.$$

We will now show that $S_{d,d}$ converges in distribution towards a normal distribution as $d \rightarrow \infty$, using a Martingale central limit theorem, see [24, Thm. 3.2, p. 58].

The set $\{S_{d,j} : 1 \leq j \leq d, d \in \mathbb{N}\}$ is a zero mean, square-integrable Martingale array, i.e. for each $d \in \mathbb{N}$ and $1 \leq j \leq d$, $S_{d,j}$ is measurable,

$$\mathbb{E}_{\pi_d}[S_{d,j}] = 0, \quad \mathbb{E}_{\pi_d}[|S_{d,j}|] < \infty, \quad \text{and} \quad \mathbb{E}_{\pi_d}[(S_{j,d})^2] < \infty.$$

For definitions, see [24, p. 1 and 53]. Define $X_{d,j} := S_{d,j} - S_{d,j-1}$. To ensure we satisfy the conditions for [24, Thm. 3.2] we must show that there exists an a.s. finite random variable η^2 such that

$$\max_{1 \leq i \leq d} |X_{d,i}| \xrightarrow{p} 0 \quad \text{as } d \rightarrow \infty, \quad (4.7)$$

$$\sum_{i=1}^d X_{d,i}^2 \xrightarrow{p} \eta^2 \quad \text{as } d \rightarrow \infty, \text{ and} \quad (4.8)$$

$$\mathbb{E}_{\pi_d} \left(\max_{1 \leq i \leq d} X_{d,i}^2 \right) \text{ is bounded in } d. \quad (4.9)$$

First consider (4.7). We have

$$|X_{d,i}| \leq |T_{1i}|(|\xi_i| + |\kappa_{d,i}|) + |T_{2i}||\nu_i| + |T_{3i}|(\xi_i^2 + \gamma_{d,i}) + |T_{4i}|(\nu_i^2 + 1) + |T_{5i}||\xi_i||\nu_i|$$

which goes to zero in probability since $\kappa_{d,i}$ and $\gamma_{d,i}$ are bounded uniformly and $|T_{ji}|$ are all $\mathcal{O}(d^{-1/2})$ as $d \rightarrow \infty$ uniformly in i .

Now consider (4.8). Define

$$\eta^2 := \sigma_{ng}^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 + (\kappa_{d,i}T_{1i} + T_{3i}(\gamma_{d,i} - 1))^2 < \infty$$

and $Y_{d,i} := dX_{d,i}^2$ so that $\bar{Y}_d = \frac{1}{d} \sum_{i=1}^d Y_{d,i} = \sum_{i=1}^d X_{d,i}^2$. Then

$$\mathbb{E}_{\tilde{\pi}_d}[Y_{d,i}] = d(T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 + (\kappa_{d,i}T_{1i} + T_{3i}(\gamma_{d,i} - 1))^2),$$

and

$$\text{Var}_{\tilde{\pi}_d}[Y_{d,i}] = d^2 \sum_{|\omega|=4} C_\omega T_{1i}^{\omega_1} T_{2i}^{\omega_2} T_{3i}^{\omega_3} T_{4i}^{\omega_4} T_{5i}^{\omega_5}$$

where ω is a multi-index with $\omega_i \geq 0$ and $|\omega| = \sum_i \omega_i = 4$, and C_ω are uniformly bounded constants. Since T_{ji} are all $\mathcal{O}(d^{-1/2})$, $\mathbb{E}_{\tilde{\pi}_d}[Y_{d,i}]$ and $\text{Var}_{\tilde{\pi}_d}[Y_{d,i}]$ are uniformly bounded.

Then by the Markov inequality, and independence of $Y_{d,i}$, for any $\epsilon > 0$,

$$\begin{aligned} \Pr(|\bar{Y}_d - \mathbb{E}_{\tilde{\pi}_d}[\bar{Y}_d]| \geq \epsilon) &\leq \frac{1}{\epsilon^2} \text{Var}_{\tilde{\pi}_d}[\bar{Y}_d] \\ &= \frac{1}{\epsilon^2 d^2} \sum_{i=1}^d \text{Var}_{\tilde{\pi}_d}[Y_{d,i}] \\ &\leq \frac{C}{\epsilon^2 d} \rightarrow 0 \quad \text{as } d \rightarrow \infty. \end{aligned}$$

Hence $\bar{Y}_d \xrightarrow{P} \eta^2$ as $d \rightarrow \infty$. This is not (4.8) yet, because it is convergence with respect to $\tilde{\pi}_d$ rather than π_d .

Since $\lim_{d \rightarrow \infty} \mathbb{E}_{\tilde{\pi}_d}[\bar{Y}_d] = \eta^2 < \infty$ and $|\bar{Y}_d| = \bar{Y}_d$, it follows that $\mathbb{E}_{\tilde{\pi}_d}[|\bar{Y}_d|]$ is uniformly bounded in d . Therefore, \bar{Y}_d is uniformly integrable and so $\bar{Y}_d \rightarrow \eta^2$ in $L^1(\tilde{\pi}_d)$ as $d \rightarrow \infty$ [33, Thm. 6.5.5 on p. 169]. Hence

$$\sum_{i=1}^d X_{d,i}^2 \xrightarrow{L^1(\tilde{\pi}_d)} \eta^2 \quad \text{as } d \rightarrow \infty.$$

From $\phi_d(x) \geq m$, the same limit holds in $L^1(\pi_d)$, which also implies convergence in probability, hence we have shown (4.8).

Condition (4.9) follows from $X_{d,i}^2 \leq \bar{Y}_d$ for $1 \leq i \leq d$, $\mathbb{E}_{\tilde{\pi}_d}[|\bar{Y}_d|]$ uniformly bounded in d , and $\phi_d(x) \geq m$.

Therefore, by the Martingale central limit theorem [24, Thm. 3.2],

$$S_{d,d} \xrightarrow{\mathcal{D}} \text{N}(0, \eta^2) \quad \text{as } d \rightarrow \infty.$$

Hence,

$$Z \xrightarrow{\mathcal{D}} \text{N}(\mu_{ng}, \sigma_{ng}^2) \quad \text{as } d \rightarrow \infty. \quad (4.10)$$

The result then follows from (4.6), (4.10) and (3.3). \blacksquare

Corollary 4.1.1 In addition to the conditions for Theorem 4.1.4, if

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + T_{3i}^2 = 0$$

then

$$\mu_{ng} = \mu \quad \text{and} \quad \sigma_{ng}^2 = \sigma^2,$$

and the expected acceptance rate for the non-Gaussian target case has the same limit as $d \rightarrow \infty$ as the Gaussian target case.

Proof. With ξ_i defined as in the proof of Theorem 4.1.4 (we only need to consider the case when matrices are diagonal), since $\mathbb{E}_{\tilde{\pi}_d}[\xi_i^2] = 1$ and $\mathbb{E}_{\tilde{\pi}_d}[\xi_i^4] = 3$ for all i and d , and since ξ_i and ξ_j are independent for $i \neq j$ under $\tilde{\pi}_d$,

$$\lim_{d \rightarrow \infty} \mathbb{E}_{\tilde{\pi}_d} \left[\left(\sum_{i=1}^d T_{1i} \xi_i + T_{3i} (\xi_i^2 - 1) \right)^2 \right] = \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + 2T_{3i}^2 = 0$$

From Jensen's inequality and $\phi_d(x) \geq m$ we have $\sum_{i=1}^d (T_{1i} \xi_i + T_{3i} (\xi_i^2 - 1)) \rightarrow 0$ in $L^1(\pi_d)$ as $d \rightarrow \infty$. Therefore, since $\kappa_{d,i} = \mathbb{E}_{\pi_d}[\xi_i]$ and $\gamma_{d,i} = \mathbb{E}_{\pi_d}[\xi_i^2]$,

$$\sum_{i=1}^d T_{1i} \kappa_{d,i} + T_{3i} (\gamma_{d,i} - 1) = \mathbb{E}_{\pi_d} \left[\sum_{i=1}^d T_{1i} \xi_i + T_{3i} (\xi_i^2 - 1) \right] \rightarrow 0 \quad \text{as } d \rightarrow \infty,$$

and $\mu_{ng} = \mu$. Also, since $\kappa_{d,i}$ and $\gamma_{d,i}$ are uniformly bounded in d and i ,

$$\sum_{i=1}^d (T_{1i} \kappa_{d,i} + T_{3i} (\gamma_{d,i} - 1))^2 \leq C \sum_{i=1}^d T_{1i}^2 + T_{3i}^2 \rightarrow 0 \quad \text{as } d \rightarrow \infty,$$

and $\sigma_{ng}^2 = \sigma^2$. ■

4.2 Expected jumpsize for non-Gaussian target

Theorem 4.2.1 Under the same conditions as Theorem 4.1.4,

$$\mathbb{E}_{\pi_d}[(q_i^T(x' - x))^2] = \left(\left(\tilde{g}_i^2 \hat{r}_i^2 + \frac{g_i}{\lambda_i^2} \right) \mathbb{E}_{\pi_d}[\alpha(x, y)] + 2 \frac{\hat{r}_i \tilde{g}_i^2 \gamma_{d,i}^{1/2}}{\lambda_i} u_{d,i} + \frac{\tilde{g}_i^2 \gamma_{d,i}}{\lambda_i^2} v_{d,i} \right) + o(t_{d,i} \lambda_i^{-1})$$

(uniformly in i) as $d \rightarrow \infty$, for some $-1 \leq u_{d,i} \leq 1$ and $0 \leq v_{d,i} \leq 1$.

Proof. As in earlier proofs, it is sufficient to prove the result in the case when all matrices are diagonal, so let A , G and Σ be diagonal matrices.

Let S_d denote the expected squared jump size in coordinate direction i , so that

$$S_d = \mathbb{E}_{\pi_d}[(x'_i - x_i)^2] = \mathbb{E}_{\pi_d}[(y_i - x_i)^2 \alpha(x, y)].$$

Also define

$$\tilde{S}_d := \mathbb{E}_{\pi_d}[(y_i - x_i)^2 \tilde{\alpha}(x, y)] \quad \text{and} \quad \tilde{S}_d^- := \mathbb{E}_{\pi_d}[(y_i - x_i)^2 \tilde{\alpha}^-(x, y)]$$

where $\tilde{\alpha}^-(x, y) = 1 \wedge \exp(\sum_{j=1, j \neq i}^d Z_{d,i})$, and $Z_{d,i}$ is the same as earlier. Recall that $\xi_i = \lambda_i(x_i - m_i)$.

We first show that $\mathbb{E}_{\pi_d}[(y_i - x_i)^4] = \mathcal{O}(t_{d,i}^2 \lambda_i^{-2})$ (uniformly in i) as $d \rightarrow \infty$. As in the proof of Theorem 3.2.1, from $y = Gx + g + \Sigma^{1/2} \nu$ where $\nu \sim \mathcal{N}(0, I)$, and $\tilde{m} = G\tilde{m} + g$ it follows that

$$y_i - x_i = -\tilde{g}_i \hat{r}_i - \frac{\tilde{g}_i}{\lambda_i} \xi_i + \frac{g_i^{1/2}}{\lambda_i} \nu_i$$

where $\nu_i \stackrel{\text{iid}}{\sim} \text{N}(0, 1)$, $\xi_i = \lambda_i(x_i - m_i)$ and $x \sim \pi_d$. Therefore,

$$(y_i - x_i)^4 = \lambda_i^{-2} \left(-\tilde{g}_i \hat{r}_i \lambda_i^{1/2} - \frac{\tilde{g}_i}{\lambda_i^{1/2}} \xi_i + \frac{\hat{r}_i^{1/4} g_i^{1/2}}{\tilde{\lambda}_i^{1/2}} \nu_i \right)^4.$$

In the proof of Theorem 4.1.4 we showed that $|\kappa_{d,i}| = |\mathbb{E}_{\pi_d}[\xi_i]|$ and $\gamma_{d,i} = \mathbb{E}_{\pi_d}[\xi_i^2]$ are uniformly bounded. Similarly, $|\mathbb{E}_{\pi_d}[\xi_i^3]|$ and $\mathbb{E}_{\pi_d}[\xi_i^4]$ are also uniformly bounded. Using these facts and $\tilde{g}_i^2 \hat{r}_i^2 \lambda_i$, $\tilde{g}_i^2 \lambda^{-1}$, $g_i \tilde{\lambda}_i^{-1}$ ($t_{d,i}$) (uniformly in i) and \hat{r}_i bounded uniformly in d and i , it follows that $\mathbb{E}_{\pi_d}[(y_i - x_i)^4] = \mathcal{O}(t_{d,i}^2 \lambda_i^{-2})$ (uniformly in i) as $d \rightarrow \infty$.

Now let us show that $S_d - \tilde{S}_d = o(t_{d,i} \lambda_i^{-1})$ as $d \rightarrow \infty$. From the Lipschitz continuity of $z \mapsto 1 \wedge \exp(z)$ and the Cauchy-Schwarz inequality,

$$\begin{aligned} S_d - \tilde{S}_d &= \mathbb{E}_{\pi_d}[(y_i - x_i)^2(\alpha(x, y) - \tilde{\alpha}(x, y))] \\ &\leq \mathbb{E}_{\pi_d}[(y_i - x_i)^4]^{1/2} \mathbb{E}_{\pi_d}[|\phi_d(x) - \phi_d(y)|^2]^{1/2} = o(t_{d,i} \lambda_i^{-1}) \quad \text{as } d \rightarrow \infty, \end{aligned}$$

by Lemma 4.1.1, since $\mathbb{E}_{\pi_d}[(y_i - x_i)^4] = \mathcal{O}(t_{d,i}^2 \lambda_i^{-2})$.

Now show that $\tilde{S}_d - \tilde{S}_d^- = o(t_{d,i} \lambda_i^{-1})$ as $d \rightarrow \infty$. Again, by the Lipschitz continuity of $z \mapsto 1 \wedge \exp(z)$ and the Cauchy-Schwarz inequality,

$$\begin{aligned} \tilde{S}_d - \tilde{S}_d^- &= \mathbb{E}_{\pi_d}[(y_i - x_i)^2(\tilde{\alpha}(x, y) - \tilde{\alpha}^-(x, y))] \\ &\leq \mathbb{E}_{\pi_d}[(y_i - x_i)^4]^{1/2} \mathbb{E}_{\pi_d}[Z_{d,i}^2]^{1/2} = o(t_{d,i} \lambda_i^{-1}) \quad \text{as } d \rightarrow \infty, \end{aligned}$$

since $\mathbb{E}_{\pi_d}[(y_i - x_i)^4] = \mathcal{O}(t_{d,i}^2 \lambda_i^{-2})$ and T_{ji} are all $\mathcal{O}(d^{-1/2})$.

Now consider \tilde{S}_d^- . Since ν_i is independent of ξ_i and $\tilde{\alpha}^-(x, y)$, $\mathbb{E}[\nu_i] = 0$ and $\mathbb{E}[\nu_i^2] = 1$,

$$\begin{aligned} \tilde{S}_d^- &= \mathbb{E}_{\pi_d}[(y_i - x_i)^2 \tilde{\alpha}^-(x, y)] \\ &= \mathbb{E}_{\pi_d} \left[\left(\tilde{g}_i^2 \hat{r}_i^2 + 2 \frac{\hat{r}_i \tilde{g}_i^2}{\lambda_i} \xi_i + \frac{\tilde{g}_i^2}{\lambda_i^2} \xi_i^2 + \frac{g_i}{\tilde{\lambda}_i^2} \right) \tilde{\alpha}^-(x, y) \right] \\ &= \left(\tilde{g}_i^2 \hat{r}_i^2 + \frac{g_i}{\tilde{\lambda}_i^2} \right) \mathbb{E}_{\pi_d}[\tilde{\alpha}^-(x, y)] + 2 \frac{\hat{r}_i \tilde{g}_i^2}{\lambda_i} \mathbb{E}_{\pi_d}[\xi_i \tilde{\alpha}^-(x, y)] + \frac{\tilde{g}_i^2}{\lambda_i^2} \mathbb{E}_{\pi_d}[\xi_i^2 \tilde{\alpha}^-(x, y)]. \end{aligned}$$

Since $\tilde{\alpha}^-(x, y) \in (0, 1]$, it follows from Jensen's inequality that

$$|\mathbb{E}_{\pi_d}[\xi_i \tilde{\alpha}^-(x, y)]| \leq \mathbb{E}_{\pi_d}[|\xi_i| \tilde{\alpha}^-(x, y)] \leq \mathbb{E}_{\pi_d}[|\xi_i|] \leq \mathbb{E}_{\pi_d}[\xi_i^2]^{1/2} = \gamma_{d,i}^{1/2}.$$

Also, $0 \leq \mathbb{E}_{\pi_d}[\xi_i^2 \tilde{\alpha}^-(x, y)] \leq \mathbb{E}_{\pi_d}[\xi_i^2] = \gamma_{d,i}$.

Finally, using $z \mapsto 1 \wedge \exp(z)$ Lipschitz, and since T_{ji} are $\mathcal{O}(d^{-1/2})$,

$$|\mathbb{E}_{\pi_d}[\tilde{\alpha}^-(x, y)] - \mathbb{E}_{\pi_d}[\tilde{\alpha}(x, y)]| \leq \mathbb{E}_{\pi_d}|Z_{d,i}| \rightarrow 0 \quad \text{as } d \rightarrow \infty.$$

Since $\tilde{g}_i^2 \hat{r}_i^2 + \frac{g_i}{\tilde{\lambda}_i^2} = \mathcal{O}(t_{d,i} \lambda_i^{-1})$ it follows that as $d \rightarrow \infty$

$$\begin{aligned} S_d &= S_d^- + o(t_{d,i} \lambda_i^{-1}) \\ &= \left(\left(\tilde{g}_i^2 \hat{r}_i^2 + \frac{g_i}{\tilde{\lambda}_i^2} \right) \mathbb{E}_{\pi_d}[\tilde{\alpha}(x, y)] + 2 \frac{\hat{r}_i \tilde{g}_i^2 \gamma_{d,i}^{1/2}}{\lambda_i} u_{d,i} + \frac{\tilde{g}_i^2 \gamma_{d,i}}{\lambda_i^2} v_{d,i} \right) + o(t_{d,i} \lambda_i^{-1}) \end{aligned}$$

for some $u_{d,i} \in [-1, 1]$ and $v_{d,i} \in [0, 1]$. The result then follows from (4.6) and $\tilde{g}_i^2 \hat{r}_i^2 + \frac{g_i}{\tilde{\lambda}_i^2} = \mathcal{O}(t_{d,i} \lambda_i^{-1})$. \blacksquare

Chapter 5

Examples

5.1 Discretized Langevin diffusion - MALA and SLA

The proposal for MALA is obtained from the Euler-Maruyama discretization of a Langevin diffusion process $\{z_t\}$ which satisfies the stochastic differential equation

$$\frac{dz_t}{dt} = \frac{1}{2} \nabla \log \pi(z_t) + \frac{dW_t}{dt},$$

where W_t is standard Brownian motion in \mathbb{R}^d . This diffusion process has the desired target distribution π as equilibrium, so one might expect a discretization of the diffusion process to almost preserve the desired target distribution. If the target is Gaussian, $N(A^{-1}b, A^{-1})$, then for current state $x \in \mathbb{R}^d$ and time step $h > 0$, the MALA proposal $y \in \mathbb{R}^d$ is defined as

$$y = (I - \frac{h}{2}A)x + \frac{h}{2}b + \sqrt{h}\xi \quad (5.1)$$

where $\xi \sim N(0, I)$. One can also use this proposal for situations where the target distribution is a change of measure from a Gaussian; in which case it is called the SLA proposal [9]. The SLA algorithm is

$$\begin{aligned} \text{Target:} & \quad \pi_d, \\ \text{Proposal:} & \quad y = (I - \frac{h}{2}A)x + \frac{h}{2}b + \sqrt{h}\xi, \quad \text{where } \xi \sim N(0, I), \end{aligned} \quad (5.2)$$

Identifying (5.1) with (1.1) and applying Corollary 2.1.1 we have the following theorem.

Theorem 5.1.1 The SLA proposal corresponds to the matrix splitting

$$\begin{aligned} M &= \frac{2}{h}(I - \frac{h}{4}A), & \mathcal{A} &= (I - \frac{h}{4}A)A, \\ N &= \frac{2}{h}(I - \frac{h}{4}A)(I - \frac{h}{2}A), & \beta &= (I - \frac{h}{4}A)b. \end{aligned}$$

Thus, the SLA proposal corresponds to a matrix splitting where M and N are functions of A and our theory applies. An important feature of this proposal is that $\hat{r}_i = 0$ for all i . This greatly simplifies the results in Theorems 3.1.2, 3.2.1, 4.1.4 and 4.2.1 and we extend existing theory ([9, Cor. 1] and the simpler results in [31, Thm. 7]) to the case where the reference Gaussian measure is allowed to have off-diagonal covariance terms. The theorem below is a special case of Theorem 5.2.2 so we omit the proof.

Theorem 5.1.2 Suppose there exist constants $c, C > 0$ and $\kappa \geq 0$ such that the eigenvalues λ_i^2 of A satisfy

$$ci^\kappa \leq \lambda_i \leq Ci^\kappa \quad \text{for } i = 1, \dots, d.$$

Also suppose that ϕ_d satisfies Assumption 4.1.2 and $r = \max\{s, s', s''\}$ satisfies Assumption 4.1.3.

If $h = l^2 d^{-1/3-2\kappa}$ for some $l > 0$ then SLA, in equilibrium, satisfies

$$\mathbb{E}[\alpha(x, y)] \rightarrow 2\Phi\left(-\frac{l^3\sqrt{\tau}}{8}\right)$$

and

$$\mathbb{E}[(x_i - x_i)^2] = 2h\Phi\left(-\frac{l^3\sqrt{\tau}}{8}\right) + o(h) \quad (5.3)$$

as $d \rightarrow \infty$ where $\tau = \lim_{d \rightarrow \infty} \frac{1}{d^{1+6\kappa}} \sum_{i=1}^d \lambda_i^6$.

Thus, the performance of SLA depends on the choice of h which is usually tuned (by tuning l) to maximise the expected jump distance. From (5.3), using $s = l\tau^{1/6}/2$, we have

$$\max_{l>0} 2l^2 d^{-1/3-2\kappa} \Phi\left(-\frac{l^3\sqrt{\tau}}{8}\right) = \max_{s>0} \frac{8d^{-1/3-2\kappa}}{\tau^{1/3}} s^2 \Phi(-s^3), \quad (5.4)$$

which is maximised at $s_0 = 0.8252$, independent of τ . Therefore, the acceptance rate that maximises expected jump distance is $2\Phi(-s_0^3) = 0.574$. This result was first stated in [30] for product target distributions, then more generally in [31, 9]. Our result is even more general because we allow the reference measure, which must be Gaussian in our case, to have off-diagonal covariance terms. In practice, h (equivalently l) is adjusted so that the acceptance rate is approximately 0.574 to maximise the expected jump size. This acceptance rate is independent of τ , so it is independent of the eigenvalues of A . However, the efficiency of SLA still depends on τ , and as τ increases the expected square jump distance will decrease by a factor $\tau^{1/3}$.

The rationale for studying the case when $d \rightarrow \infty$ is that it is a good approximation for cases when d is ‘large’ and finite (see eg. [30, Fig. 1] or [31]). However, the results above suggest that we should take $h \rightarrow 0$ as $d \rightarrow \infty$, to achieve at best an expected jump size that also tends towards 0 as $d \rightarrow \infty$. The only good point about these results is that SLA has superior asymptotic performance over RWM (see [9, Thms. 1-4 and Cor. 1] and [30, Fig. 1]).

To understand the convergence of SLA to equilibrium (burn in) we would like to know the ‘spectral gap’ or second largest eigenvalue of the transition kernel, as this determines the rate of convergence. As far as we are aware this is an open problem.

We can also use Theorem 5.1.1 to analyse the unadjusted Langevin algorithm (ULA) in [32] in the case when the target is Gaussian, in which case ULA is simply the proposal chain from (5.1) without the accept/reject step in the MH algorithm. From Theorem 5.1.1 we see that it does not converge to the correct target distribution since $\mathcal{A} \neq A$. Instead of converging to $N(A^{-1}b, A^{-1})$, ULA converges to $N(A^{-1}b, \mathcal{A}^{-1})$. Indeed, the authors of [32] note that ULA has poor convergence properties. We see here the reason why it converges to the wrong target distribution, and from [17, 19] we know its convergence rate to the incorrect target distribution depends on the spectral radius of $G = I - \frac{h}{2}A$, which is close to 1 when h is small.

Despite possibly having slow convergence per iteration, the SLA proposal is cheap to compute. Since $G = I - \frac{h}{2}A$, we only require a single matrix-vector multiplication with A for each proposal and since $\Sigma = hI$, an i.i.d. sample from $N(0, \Sigma)$ at each iteration is also cheap to compute.

θ	V	Method
0	I	SLA, MALA and ULA [32, 9]. SLA=MALA for Gaussian target distributions.
0	A^{-1}	Used in [28]. With a change of variables $x \leftrightarrow Q^T V^{-1/2} x$, it is the Preconditioned Simplified Langevin Algorithm (P-SLA) [9].
$\in [0, 1]$	I	θ -SLA [9].
$\frac{1}{2}$	I	CN [15].
$\frac{1}{2}$	A^{-1}	pCN [15].

Table 5.1. Different choices of θ and V in (5.6) lead to different proposals for the MH algorithm (5.7). Other choices are possible.

5.2 Discretized Langevin diffusion - more general algorithms

It is possible to generalise MALA and SLA by ‘preconditioning’ the Langevin diffusion process and using a discretization scheme that is not Euler-Maruyama. For symmetric positive definite matrix $V \in \mathbb{R}^{d \times d}$ (the ‘preconditioner’) consider a Langevin process $\{z_t\}$ satisfying the stochastic differential equation

$$\frac{dz_t}{dt} = \frac{1}{2} V \nabla \log \pi(z_t) + \frac{dW_t}{dt} \quad (5.5)$$

where W_t is Brownian motion in \mathbb{R}^d with covariance V . This diffusion process also preserves π . For $\theta \in [0, 1]$, time step $h > 0$ and current state $x \in \mathbb{R}^d$ define a proposal $y \in \mathbb{R}^d$ by discretizing (5.5) as

$$y - x = \frac{h}{2} V \nabla \log \pi(\theta y + (1 - \theta)x) + \sqrt{h} \nu$$

where $\nu \sim N(0, V)$. When the target is Gaussian, $N(A^{-1}b, A^{-1})$, this can be rewritten as

$$y = (I + \frac{\theta h}{2} V A)^{-1} \left[(I - \frac{(1-\theta)h}{2} V A)x + \frac{h}{2} V b + (hV)^{1/2} \xi \right] \quad (5.6)$$

where $\xi \sim N(0, I)$, which is an AR(1) proposal. Thus, we define a MH algorithm by

$$\begin{aligned} &\text{Target: } \pi_d, \\ &\text{Proposal: } y = (I + \frac{\theta h}{2} V A)^{-1} \left[(I - \frac{(1-\theta)h}{2} V A)x + \frac{h}{2} V b + (hV)^{1/2} \xi \right], \text{ for } \xi \sim N(0, I). \end{aligned} \quad (5.7)$$

Different choices of θ and V give different proposals. For example, SLA has $\theta = 0$ and $V = I$, and pCN corresponds to $\theta = \frac{1}{2}$ and $V = A^{-1}$. Table 5.1 describes several more examples.

Applying Corollary 2.1.1 we can prove the following theorem.

Theorem 5.2.1 The general Langevin proposal (5.6) corresponds to the matrix splitting

$$\begin{aligned} M &= \frac{2}{h} V^{-1/2} W (I + \frac{\theta h}{2} B) V^{-1/2}, & \mathcal{A} &= V^{-1/2} W B V^{-1/2} = \tilde{W} A, \\ N &= \frac{2}{h} V^{-1/2} W (I - \frac{(1-\theta)h}{2} B) V^{-1/2}, & \beta &= V^{-1/2} W V^{1/2} b = \tilde{W} b, \end{aligned}$$

where $B = V^{1/2} A V^{1/2}$, $W = I + (\theta - \frac{1}{2}) \frac{h}{2} B$ and $\tilde{W} = I + (\theta - \frac{1}{2}) \frac{h}{2} A V$.

Proof. With iteration matrix $G = (I + \frac{\theta h}{2}VA)^{-1}(I - \frac{(1-\theta)h}{2}VA)$, vector $g = \frac{h}{2}(I + \frac{\theta h}{2}VA)^{-1}Vb$ and matrix $\Sigma = (I + \frac{\theta h}{2}VA)^{-1}(hV)(I + \frac{\theta h}{2}AV)^{-1}$, then (5.6) is the same as (1.1). To apply Corollary 2.1.1 we first check that $G\Sigma$ is symmetric. We have

$$G = V^{1/2}(I + \frac{\theta h}{2}B)^{-1}(I - \frac{(1-\theta)h}{2}B)V^{-1/2} \quad \text{and} \quad \Sigma = hV^{1/2}(I + \frac{\theta h}{2}B)^{-2}V^{1/2}$$

so that

$$G\Sigma = hV^{1/2}(I + \frac{\theta h}{2}B)^{-1}(I - \frac{(1-\theta)h}{2}B)(I + \frac{\theta h}{2}B)^{-2}V^{1/2}$$

which is symmetric since V and B are symmetric. Applying Corollary 2.1.1 then yields the result. \blacksquare

Because of Theorem 5.2.1, the proposal chain for (5.6) converges to $N(A^{-1}b, (\tilde{W}A)^{-1})$, and if $\theta \neq \frac{1}{2}$, then $\tilde{W} \neq I$, $\mathcal{A} \neq A$, and the target reference and proposal limit distributions disagree.

If $\theta = \frac{1}{2}$, then the proposal limit and target reference distributions are the same. If, in addition, the target is Gaussian, then the MH accept/reject step is redundant and we can use [17, 19, 18] to analyse and accelerate the proposal chain generated by (5.6).

To evaluate the performance of the MH algorithm (5.7) when $\mathcal{A} \neq A$ we would like to be able to apply Theorems 4.1.4 and 4.2.1, but we see in Theorem 5.2.1 that the splitting matrices are functions of B and V (not A) so we cannot directly apply our new theory. However, a change of coordinates will fix this! The following lemma is a result of simple algebra and Theorem 5.2.1.

Lemma 5.2.1 Under the change of coordinates

$$x \leftrightarrow V^{-1/2}x$$

the MH algorithm (5.7) is transformed to the MH algorithm defined by

$$\begin{aligned} \text{Target: } \pi_d \quad \text{where } \frac{d\pi_d}{d\tilde{\pi}_d}(x) &= \exp(-\phi_d(V^{1/2}x)) \quad \text{and } \tilde{\pi}_d \text{ is } N(B^{-1}V^{1/2}b, B^{-1}), \\ \text{Proposal: } (I + \frac{\theta h}{2}B)y &= (I - \frac{(1-\theta)h}{2}B)x + \frac{h}{2}V^{1/2}b + h^{1/2}\xi, \quad \text{where } \xi \sim N(0, I), \end{aligned} \quad (5.8)$$

and $B = V^{1/2}AV^{1/2}$. Moreover, the proposal in (5.8) corresponds to the matrix splitting $B = M - N$ where

$$\begin{aligned} M &= \frac{2}{h}W(I + \frac{\theta h}{2}B), & \mathcal{A} &= WB, \\ N &= \frac{2}{h}W(I - \frac{(1-\theta)h}{2}B), & \beta &= WV^{1/2}b, \end{aligned}$$

and $W = I + (\theta - \frac{1}{2})\frac{h}{2}B$.

Thus, we have transformed the MH algorithm (5.7) to a MH algorithm where the splitting matrices are functions of the target reference precision matrix, and we can apply Theorems 4.1.4 and 4.2.1 to (5.8) to find the expected acceptance rate and expected jump size of (5.7). Note that we never compute the Markov chain for (5.8), we only use it to determine the convergence properties of (5.7) since they are identical.

Theorem 5.2.2 Suppose there are constants $c, C > 0$ and $\kappa \geq 0$ such that the eigenvalues λ_i^2 of $B = V^{1/2}AV^{1/2}$ (equivalently, λ_i^2 are eigenvalues of VA) satisfy

$$ci^\kappa \leq \lambda_i \leq Ci^\kappa \quad \text{for } i = 1, \dots, d.$$

Also suppose that $\psi_d(x) := \phi_d(V^{1/2}x)$ satisfies Assumption 4.1.2 (with $|\cdot|_s = |B^s \cdot|$ for $s \in \mathbb{R}$) and $r = \max\{s, s', s''\}$ satisfies Assumption 4.1.3.

If $h = l^2 d^{-1/3-2\kappa}$ for $l > 0$ and $\tau = \lim_{d \rightarrow \infty} \frac{1}{d^{6\kappa+1}} \sum_{i=1}^d \lambda_i^6$ then MH algorithm (5.7), in equilibrium, satisfies

$$\mathbb{E}[\alpha(x, y)] \rightarrow 2\Phi\left(-\frac{l^3|\theta - \frac{1}{2}|\sqrt{\tau}}{4}\right) \quad (5.9)$$

and for normalised eigenvector q_i of B corresponding to λ_i^2 ,

$$\mathbb{E}[|q_i^T V^{-1/2}(x' - x)|^2] = 2h\Phi\left(-\frac{l^3|\theta - \frac{1}{2}|\sqrt{\tau}}{4}\right) + o(h) \quad (5.10)$$

as $d \rightarrow \infty$.

The proof of Theorem 5.2.2 is in the Appendix.

For efficiency, as well as considering the expected squared jump size, we must also consider the computing cost of the proposal (5.6), which requires the action of $(I + \frac{\theta h}{2}VA)^{-1}$ and an independent sample from $N(0, V)$, as well as the actions of V and A multiplying a vector.

We see that V plays a similar role to a preconditioner in solving a linear system. For a linear system we choose V to be cheap to compute matrix multiplication and to minimise the condition number of VA . For the MH algorithm (5.7) we choose it so that multiplying with V and sampling from $N(0, V)$ are cheap to compute and to minimise τ . In both cases V is chosen to ‘control’ the eigenvalues of VA .

Although SLA and more general discretizations of Langevin diffusion have been successfully analyzed in [9], all of these results are stated for target distributions that are a change of measure from product measures. Theorem 5.2.2 extends their theory (in particular [9, Cor. 1]) to the case where the reference measure $\tilde{\pi}_d$ may be Gaussian with off-diagonal covariance terms and $\theta \in [0, 1]$. See also [31, Thm. 7].

With $\theta = 0$ and $V = I$, SLA is very cheap to compute because we only have to invert the identity matrix and sample from $N(0, I)$ at each iteration (as well as multiply by A). Alternatively, pCN, with $\theta = \frac{1}{2}$ and $V = A^{-1}$, requires a sample from $N(0, A^{-1})$ which may be computationally expensive, particularly in high dimensions.

5.3 L -step methods

Given an AR(1) proposal of the form (1.1), we can form a new AR(1) proposal by taking L steps of the original proposal before performing the MH accept/reject step. This may be advantageous when the cost of evaluating ϕ_d is significant. The L -step proposal is computed by iterating

$$y^{(l)} = Gy^{(l-1)} + g + \nu^{(l)} \quad \text{for } l = 1, \dots, L,$$

where $\nu^{(l)}$ is an i.i.d. draw from $N(0, \Sigma)$ and $y^{(0)} = x$. This yields a new proposal in the form of (1.1),

$$y = G_L x + g_L + \nu_L, \quad \text{with } \nu_L \sim N(0, \Sigma_L) \quad (5.11)$$

where $G_L = G^L$, $g_L = (I - G)^{-1}(I - G^L)g$ and $\Sigma_L = \sum_{l=0}^{L-1} G^l \Sigma (G^T)^l$. Hence, the eigenvalues of G_L are G_i^L , and if $G_i < 1$ then the L -step proposal chain will converge to the same limit as the 1-step proposal chain (i.e. $\mathcal{A}_L = \mathcal{A}$ and $\beta_L = \beta$).

We can reduce the computational cost of evaluating the acceptance ratio for the L -step proposal using the surrogate transition method [25, p.194]. The proof of the following lemma is in the Appendix.

Lemma 5.3.1 The L -step acceptance probability satisfies

$$\alpha(x, y) = 1 \wedge \frac{\pi_d(y)q_L(y, x)}{\pi_d(x)q_L(x, y)} = 1 \wedge \frac{\pi_d(y)\pi^*(x)}{\pi_d(x)\pi^*(y)}$$

where $q_L(x, dy) = q_L(x, y)dy$ is the transition kernel for the L -step proposal y given x from (5.11) and $\pi^*(x) \propto \exp(-\frac{1}{2}x^T Ax + \beta^T x)$.

The computational cost of the L -step proposal is L times the cost of the original proposal, but the expected squared jump size for the L -step method is, in general, not L times the original. For example, let us consider L -step SLA, where $G = (I - \frac{h}{2}A)$, $\beta = \frac{h}{2}b$ and $\Sigma = hI$, and the proposal is given by (5.11). The proof of the following theorem is in the Appendix.

Theorem 5.3.1 Suppose there exist constants $c, C > 0$ and $\kappa \geq 0$ such that the eigenvalues λ_i^2 of A satisfy

$$ci^\kappa \leq \lambda_i \leq Ci^\kappa \quad \text{for } i = 1, \dots, d.$$

Also suppose that ϕ_d satisfies Assumption 4.1.2 and $r = \max\{s, s', s''\}$ satisfies Assumption 4.1.3.

If $h = l^2 d^{-1/3-2\kappa}$ for some $l > 0$ then L -step SLA, in equilibrium, satisfies

$$\mathbb{E}[\alpha(x, y)] \rightarrow 2\Phi\left(-\frac{l^3\sqrt{L\tau}}{8}\right) \quad (5.12)$$

and

$$\mathbb{E}[(x'_i - x_i)^2] = 2Lh\Phi\left(-\frac{l^3\sqrt{L\tau}}{8}\right) + o(h) \quad (5.13)$$

as $d \rightarrow \infty$ where $\tau = \lim_{d \rightarrow \infty} \frac{1}{d^{1+6\kappa}} \sum_{i=1}^d \lambda_i^6$.

To maximise the performance of L -step SLA we then tune l to maximise the expected jump size. From (5.13), using $s = l(L\tau)^{1/6}/2$, we have

$$\max_{l>0} 2Ll^2 d^{-1/3-2\kappa} \Phi\left(-\frac{l^3\sqrt{L\tau}}{8}\right) = \max_{s>0} 8L^{2/3} d^{-1/3-2\kappa} \tau^{-1/3} s^2 \Phi(-s^3), \quad (5.14)$$

which is maximised at $s_0 = 0.8252$. Therefore, the expected jump size of L -step SLA is maximised when the acceptance rate is $2\Phi(-s_0^3) = 0.574$, which is the same as SLA, but this corresponds to an expected jump size that is only $L^{2/3}$ times larger than the jump size for SLA (compare (5.14) and (5.4)) in the limit when $d \rightarrow \infty$.

To compare the efficiency L -step SLA for varying L we must also consider the computational cost of the method. For example, suppose that matrix-vector products with A cost 1 unit of CPU time, inner products and drawing independent samples from $N(0, I)$ are essentially free, and evaluating ϕ_d costs t units of CPU time. From Lemma 5.3.1, we can simplify the acceptance ratio for L -step SLA to

$$\alpha(x, y) = 1 \wedge \exp\left(\frac{h}{4}(|Ax|^2 - |Ay|^2) - \frac{h}{2}b^T(Ax - Ay) + \phi_d(x) - \phi_d(y)\right)$$




Figure 5.1. Efficiency of the L -step SLA method for varying number of steps L and varying computing cost for evaluating ϕ_d . Filled markers correspond to maximum efficiency.

so L -step SLA uses L matrix vector products with A per proposal and an additional matrix-vector product and two evaluations of ϕ_d in the acceptance ratio. If the proposal is accepted then we can reuse some of the calculations in the acceptance ratio, but if it is rejected then a matrix-vector product and an evaluation of ϕ_d are wasted. The average cost of an L -step SLA iteration is then

$$L + t + (1 - \alpha)(1 + t) = 1.426 + 0.426t + L$$

units of CPU time, assuming that we have tuned L -step SLA so that the acceptance rate is 0.574. Also let 1 unit of jump size be the expected jump size of 1-step SLA, then L -step SLA has an expected jump size of $L^{2/3}$ units, and the ‘efficiency’ of L -step SLA is calculated as jump size divided by computing cost,

$$\frac{L^{2/3}}{1.426 + 0.426t + L},$$

which is maximised at $L = 2(1.426 + 0.426t)$. Our conclusion is that SLA can be improved by using L -step SLA with $L > 1$, and the optimal value of L depends on the cost of evaluating ϕ_d . If $t = 0$, then $L = 3$ is optimal. Figure 5.1 shows the efficiency of L -step SLA for other values of t .

This analysis can be repeated for other L -step algorithms.

5.4 Hybrid Monte Carlo

Another type of AR(1) proposal, that fits our theory when π_d is Gaussian, are proposals from the Hybrid (or Hamiltonian) Monte Carlo algorithm (HMC), see e.g. [16, 6, 26]. For this section, suppose that the target π_d is the Gaussian $N(A^{-1}b, A^{-1})$.

HMC treats the current state $x \in \mathbb{R}^d$ as the initial position of a particle, the initial momentum $p \in \mathbb{R}^d$ of the particle is chosen independently at random, and then the motion of the particle is evolved according to a Hamiltonian system for a fixed amount of time. The final position of the particle is the proposal. Instead of solving the Hamiltonian system exactly, the evolution of the particle is approximated using a reversible, symplectic numerical integrator. For example, the leap-frog method (also called the Stormer-Verlet method) is an integrator that preserves a modified Hamiltonian, see e.g. [23]. Hence, the proposal y , for target $N(A^{-1}b, A^{-1})$, is computed as follows; let $V \in \mathbb{R}^{d \times d}$ be a symmetric positive definite matrix and define a Hamiltonian function $H : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$H(q, p) := \frac{1}{2}p^T V p + \frac{1}{2}q^T A q - b^T q. \quad (5.15)$$

Given a time step $h > 0$, a number of steps $L \in \mathbb{N}$, and current state $x \in \mathbb{R}^d$, define $q_0 := x$, and sample $p_0 \sim N(0, V^{-1})$. Then for $l = 0, \dots, L-1$ compute

$$\begin{aligned} p_{l+1/2} &= p_l - \frac{h}{2}(Aq_l - b), \\ q_{l+1} &= q_l + hVp_{l+1/2}, \\ p_{l+1} &= p_{l+1/2} - \frac{h}{2}(Aq_{l+1} - b). \end{aligned}$$

The proposal is then defined as $y := q_L$. In matrix form we have

$$\begin{bmatrix} q_{l+1} \\ p_{l+1} \end{bmatrix} = K \begin{bmatrix} q_l \\ p_l \end{bmatrix} + J \begin{bmatrix} 0 \\ \frac{h}{2}b \end{bmatrix}$$

where $K, J \in \mathbb{R}^{2d \times 2d}$ are defined as

$$K = \begin{bmatrix} I & 0 \\ -\frac{h}{2}A & I \end{bmatrix} \begin{bmatrix} I & hV \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ -\frac{h}{2}A & I \end{bmatrix} = \begin{bmatrix} I - \frac{h^2}{2}VA & hV \\ -hA + \frac{h^3}{4}AVA & I - \frac{h^2}{2}AV \end{bmatrix}$$

and

$$J = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} + \begin{bmatrix} I & 0 \\ -\frac{h}{2}A & I \end{bmatrix} \begin{bmatrix} I & hV \\ 0 & I \end{bmatrix} = \begin{bmatrix} 2I & hV \\ -\frac{h}{2}A & 2I - \frac{h^2}{2}AV \end{bmatrix}.$$

Hence, y is given by

$$\begin{bmatrix} y \\ p_L \end{bmatrix} = K^L \begin{bmatrix} x \\ \xi \end{bmatrix} + \sum_{l=0}^{L-1} K^l J \begin{bmatrix} 0 \\ \frac{h}{2}b \end{bmatrix} \quad \text{where } \xi \sim N(0, V), \quad (5.16)$$

or equivalently,

$$y = (K^L)_{11}x + \left(S J \begin{bmatrix} 0 \\ \frac{h}{2}b \end{bmatrix} \right)_1 + (K^L)_{12}\xi, \quad \text{where } \xi \sim N(0, V^{-1}), \quad (5.17)$$

and $(K^L)_{ij}$ is the ij block (of size $d \times d$) of K^L , $S = (I - K)^{-1}(I - K^L)$ and $(\cdot)_1$ are the first d entries of the vector (\cdot) .

In the case of only one time step of the leap-frog integrator ($L = 1$) then HMC is MALA [6]. Hence, we immediately know that the HMC proposal with $L = 1$ is an AR(1) proposal where the proposal limit and target distributions are not the same, and the expected acceptance rate and jump size are given by Theorem 5.1.2. The case for $L > 1$ is more complicated, but (5.17) is still an AR(1) proposal that can be expressed as a matrix splitting using (1.5). The proofs of the following two results are in the Appendix.

Theorem 5.4.1 The HMC proposal (5.17) corresponds to the matrix splitting

$$\begin{aligned} M &= \Sigma^{-1}(I + (K^L)_{11}), & \mathcal{A} &= \Sigma^{-1}(I - (K^L)_{11}^2), \\ N &= \Sigma^{-1}(I + (K^L)_{11})(K^L)_{11}, & \beta &= \Sigma^{-1}(I + (K^L)_{11}) \left((\mathcal{S}\mathcal{J} \begin{bmatrix} 0 \\ \frac{h}{2}b \end{bmatrix}) \right)_1, \end{aligned}$$

where $\Sigma = (K^L)_{12}V^{-1}(K^L)_{12}^T$.

Corollary 5.4.1 The matrix splitting from HMC satisfies $\mathcal{A}^{-1}\beta = A^{-1}b$.

These results imply that the proposal chain for HMC converges to $N(A^{-1}b, \mathcal{A}^{-1})$ where $\mathcal{A} \neq A$, rather than the desired target $N(A^{-1}b, A^{-1})$, so the MH accept/reject step is necessary even when the target is Gaussian.

For the analysis of HMC we require the eigenvalues of the iteration matrix. A proof of the following result is in the Appendix.

Theorem 5.4.2 Let λ_i^2 be eigenvalues of $B = V^{1/2}AV^{1/2}$ or VA (these matrices have the same eigenvalues). Then the iteration matrix $G = (K^L)_{11}$ for the HMC proposal has eigenvalues

$$G_i = \cos(L\theta_i)$$

where $\theta_i = -\cos^{-1}(1 - \frac{h^2}{2}\lambda_i^2)$.

From this theorem we see how the eigenvalues of the iteration matrix depend on V , the number of time steps L , and the time step h . Again we refer to V as a preconditioner (as in [7]) because it plays a similar role to a preconditioner for solving linear systems of equations. Alternatively, V may be referred to as a mass matrix since p in the Hamiltonian (5.15) is momentum and H is energy.

To complete our analysis of HMC we restrict our attention to the case when $d \rightarrow \infty$ and try to apply Theorems 3.1.2 and 3.2.1. These theorems require that the splitting matrices are functions of the target precision matrix. A simple change of coordinates achieves this.

Theorem 5.4.3 Under the change of coordinates

$$\begin{bmatrix} x \\ p \end{bmatrix} \leftrightarrow \mathcal{V}^{-1} \begin{bmatrix} x \\ p \end{bmatrix}, \quad \text{where } \mathcal{V} = \begin{bmatrix} V^{1/2} & 0 \\ 0 & V^{-1/2} \end{bmatrix} \in \mathbb{R}^{2d \times 2d},$$

the Hamiltonian (5.15) and HMC with target $N(A^{-1}b, A^{-1})$ and proposal (5.17) are transformed to a Hamiltonian, and MH algorithm defined by

$$\begin{aligned} \text{Hamiltonian: } & \mathcal{H}(x, p) := \frac{1}{2}p^T p + \frac{1}{2}x^T Bx - (V^{1/2}b)^T x, \\ \text{Target: } & N(B^{-1}V^{1/2}b, B^{-1}), \\ \text{Proposal: } & y = (\mathcal{K}^L)_{11}x + \left(\mathcal{S}\mathcal{J} \begin{bmatrix} 0 \\ \frac{h}{2}V^{1/2}b \end{bmatrix} \right)_1 + (\mathcal{K}^L)_{12}\xi, \text{ for } \xi \sim N(0, I), \end{aligned} \tag{5.18}$$

where $B = V^{1/2}AV^{1/2}$, $\mathcal{S} = (I - \mathcal{K})^{-1}(I - \mathcal{K}^L)$,

$$\mathcal{K} = \begin{bmatrix} I - \frac{h^2}{2}B & hI \\ -hB + \frac{h^3}{4}B^2 & I - \frac{h^2}{2}B \end{bmatrix} \quad \text{and} \quad \mathcal{J} = \begin{bmatrix} 2I & hI \\ -\frac{h}{2}B & 2I - \frac{h^2}{2}B \end{bmatrix}.$$

Moreover, the proposal in (5.18) corresponds to the matrix splitting $\mathcal{A} = M - N$ with

$$\begin{aligned} M &= (\mathcal{K}^L)_{12}^{-2}(I + (\mathcal{K}^L)_{11}), & \mathcal{A} &= (\mathcal{K}^L)_{12}^{-2}(I - (\mathcal{K}^L)_{11}^2), \\ N &= (\mathcal{K}^L)_{12}^{-2}(I + (\mathcal{K}^L)_{11})(\mathcal{K}^L)_{11}, & \beta &= (\mathcal{K}^L)_{12}^{-2}(I + (\mathcal{K}^L)_{11}) \left(\mathcal{S}\mathcal{J} \begin{bmatrix} 0 \\ \frac{h}{2}V^{1/2}b \end{bmatrix} \right)_1. \end{aligned}$$

Proof. Use $K = \mathcal{V}\mathcal{K}\mathcal{V}^{-1}$ and $J = \mathcal{V}\mathcal{J}\mathcal{V}^{-1}$. ■

Similar coordinate transformations are used in classical mechanics [1, p. 103], see also [13].

MH algorithm (5.18) has splitting matrices that are functions of the target precision matrix, so we can apply Theorems 3.1.2 and 3.2.1 to (5.18) to reveal information about the performance of the original HMC algorithm. A proof of the following result is in the Appendix.

Theorem 5.4.4 Suppose there are constants $c, C > 0$ and $\kappa \geq 0$ such that the eigenvalues of $B = V^{1/2}AV^{1/2}$ (equivalently, VA or AV) satisfy

$$ci^\kappa \leq \lambda_i \leq Ci^\kappa \quad \text{for } i = 1, \dots, d.$$

If $h = ld^{-1/4-\kappa}$ for $l > 0$, and $L = \lfloor \frac{T}{h} \rfloor$ for fixed T , then the HMC algorithm (with proposal (5.17) and target $N(A^{-1}b, A^{-1})$), in equilibrium, satisfies

$$\mathbb{E}[\alpha(x, y)] \rightarrow a(l) := 2\Phi\left(-\frac{l^2\sqrt{\tau}}{8}\right) \quad (5.19)$$

where $\tau = \lim_{d \rightarrow \infty} \frac{1}{d^{1+4\kappa}} \sum_{i=1}^d \lambda_i^4 \sin^2(\lambda_i T')$ and for eigenvector q_i of B corresponding to λ_i^2 ,

$$\mathbb{E}[|q_i^T V^{-1/2}(x' - x)|^2] = 2 \frac{1 - \cos(\lambda_i T')}{\lambda_i^2} a(l) + o\left(\frac{1 - \cos(\lambda_i T')}{\lambda_i^2}\right) \quad (5.20)$$

as $d \rightarrow \infty$, where $T' = Lh$.

The above result is an extension to the results in [7, 6] since their results only cover the situation when the target distribution has diagonal covariance, and $\kappa > 1/2$ or $\kappa = 0$.

Note that in the above theory if we take $h = ld^{-r}$ for some $r \neq \frac{1}{4} + \kappa$ then we would find that the expected acceptance rate would either tend to 0 or 1. As is well known for Metropolis-Hastings algorithms, the optimal acceptance rate usually lies somewhere between 0 and 1 and taking $h = ld^{-1/4-\kappa}$ is the correct scaling of h to the dimension to achieve this. However, even with this scaling we are still free to tune h by varying l and from the above result we can derive the optimal acceptance rate for HMC for a wider class of problems than previously studied in [6, 7].

For the efficiency of HMC, as well as considering the expected squared jump size of the chain, we must also consider the compute time per proposal which is proportional to $L = \lfloor \frac{T}{h} \rfloor$ and depends on d . Therefore, to maximise efficiency of HMC we must maximise the expected jump

size divided by the compute time for a proposal. With all other quantities held constant, this corresponds to varying l to maximise

$$la(l) = C\sqrt{s}\Phi(s)$$

where $s = \frac{l^2\sqrt{\tau}}{8}$ and C is a constant, which is maximised at $s_0 = 0.4250$, which corresponds to an expected acceptance rate of $2\Phi(s_0) = 0.651$. This is the same acceptance rate found in [6] where they considered target distributions that are product distributions with λ_i constant for all i .

Our theory goes further than earlier results in [6] and [7] for HMC because we also provide guidance on how to choose the other parameters in HMC: V and T . From Theorem 5.4.4 we see that we should choose V in a similar way to how we would choose a preconditioner for solving a linear system of equations. We should choose V to control the spread of eigenvalues of VA so that κ is as small as possible and τ is minimised, and we should choose V so that the action of matrix multiplication with V and sampling from $N(0, V^{-1})$ are cheap to compute. This result is touched on in [7] where they suggest taking $V = A^{-1}$, the perfect preconditioner, but if it is possible to sample from $N(0, A)$ then other sampling algorithms may be more efficient than HMC, particularly when the target distribution is $N(A^{-1}b, A^{-1})$.

Our theory also shows how to choose T to maximise efficiency. After tuning h to achieve an acceptance rate of 0.651 the expected squared jump size satisfies

$$\mathbb{E}[|q_i^T V^{-1/2}(x' - x)|^2] \rightarrow 1.302 \frac{1 - \cos(\lambda_i T')}{\lambda_i^2} \quad \text{as } d \rightarrow \infty.$$

By first deciding which i correspond to directions we need to consider for our statistic of interest (this will depend on the eigenvectors q_i of AV) we can then choose T to maximise $1 - \cos(\lambda_i T')$ for those i . Thus, how we choose T depends on the problem, V and our statistic of interest. In the special case when λ_i are equal then we should choose $T = \frac{\pi}{\lambda_i}$.

The theory presented here is for the leap-frog numerical integrator applied to the Hamiltonian system. Higher order integrators are also suggested in [6] and alternative numerical integrators based on splitting methods (in the ODEs context) are suggested in [13] that minimize the Hamiltonian error after L steps of the integrator. It may be possible to evaluate these other methods by first expressing them as a an AR(1) proposal and writing down the corresponding matrix splitting, then applying Theorems 3.1.2 and 3.2.1 after a change of variables; as we have done for the leap-frog integrator.

For non-Gaussian target distributions, we cannot apply Theorems 4.1.4 and 4.2.1 to this HMC algorithm (with $L = \lfloor \frac{T}{h} \rfloor$), since $\mathbb{E}_{\pi_d}[|y - x|^{2q}] \rightarrow 0$ as $d \rightarrow \infty$, since $\tilde{g}_i^2 \lambda_i^{-1}$, $g_i \tilde{\lambda}_i^{-1} \rightarrow 0$ as $d \rightarrow \infty$, see the proof of Theorem 5.4.4 in the Appendix.

Chapter 6

Concluding remarks

Until now, each MH algorithm with an AR(1) proposal has required its own analysis, e.g. RWM, MALA, SLA, pCN and HMC. In this article we have designed a unifying theory that encompasses all of these AR(1) proposals (except RWM) and other general AR(1) proposals where G and Σ are functions of A , for the case where the target distribution is a change of measure from a Gaussian reference measure.

The main analysis tool we used is matrix splitting. By writing an AR(1) proposal in terms of a matrix splitting, and requiring that the splitting matrices are functions of the target reference precision matrix A , then a simple change of variables diagonalises both the target reference A and the matrices in the proposal; G and Σ . A consequence of this fact is that it is sufficient to analyse MH algorithms where the proposal and target reference are defined by diagonal matrices. Essentially, we reduce the general case back to MH algorithms where the target reference measure is a product distribution, for which existing analysis of MALA, SLA and HMC can be extended to general AR(1) proposals.

In particular, we wrote down the obvious extension of results for Langevin proposals in [9] to the case where the target reference measure is Gaussian with non-diagonal covariance and $\theta \in [0, 1]$. For HMC, we extended results in [6] to the case when $\kappa \geq 0$ from $\kappa = 0$, we derived a new formula for the eigenvalues of the iteration matrix of the HMC proposal, and we provided criteria on how to choose T and V for HMC (previous analysis only said to adjust h until the acceptance rate is 0.651).

We also analysed a variation of the SLA algorithm where L steps of the proposal are taken before the accept/reject step. We simplified the evaluation of the acceptance probability using the surrogate transition method, and we found that, in the circumstances considered, it is optimal to take $L > 1$ steps of SLA.

The analysis presented here requires that the splitting matrices are functions of the target reference precision matrix. In high dimensions this is a natural assumption to make because factorizing A may be computationally infeasible.

Designing proposals for the MH algorithm to achieve efficient MCMC methods is a challenge, particularly for non-Gaussian target distributions, and the job is made harder by the difficulty we have in analysing the convergence properties of MH algorithms. By focusing on AR(1) proposals in high dimension we have proven new theoretical results that provide us with criteria for evaluating and constructing new AR(1) proposals for efficient MH algorithms.

Designing an efficient MH algorithm with an AR(1) proposal is often a balancing act between minimising the integrated autocorrelation time (we use maximising expected jump size as a proxy for this) and minimising compute time for each iteration of the chain. If the proposal limit and

target distributions are Gaussian and identical then it follows from the theory in [17, 19, 18] that to construct an efficient AR(1) process we should try to satisfy the following conditions:

1. The spectral radius of G should be as small as possible.
2. Computing an iteration of the stochastic AR(1) process should be as cheap. This means that the action of G and independent sampling from $N(0, \Sigma)$ should be cheap to compute.

If the target distribution is Gaussian and different from the proposal distribution then Theorems 3.1.2 and 3.2.1 suggest that, in addition, we should try to satisfy:

3. The difference between the target and proposal limit distributions should be as small as possible in the sense that the difference in means should be small, and the relative difference in precision matrix eigenvalues should be small.

If the target distribution is non-Gaussian, then Theorems 4.1.4 and 4.2.1 suggest how we should try to satisfy:

3. The difference between the target reference and proposal limit distributions should be as small as possible in the sense that the difference in means should be small, and the relative difference in precision matrix eigenvalues should be small.

In particular examples we can quantify these conditions using our theory. For example, for proposals based on discretized generalised Langevin diffusion, Theorem 5.2.2 shows us how the choice of symmetric positive definite matrix V effects efficiency as it effects squared jump size in four ways. Whilst choosing V to maximise the limit in (5.10) (by minimising κ and τ) we should balance this against the scaling and direction that V induces on $E[q_i^T V^{-1/2} (x' - x)^2]$ through q_i and $V^{-1/2}$ on the left-hand side of (5.10).

Another example, pCN, satisfies conditions 1 and 3 above, but not necessarily condition 2. In particular, G is the diagonal matrix with entries all $\frac{1-h/4}{1+h/4}$ on the diagonal, and $\mathcal{A} = A$ and $\beta = b$. However, each proposal for pCN requires an independent sample from $N(0, A^{-1})$, which may be infeasible in high dimensions. In the special case when A is diagonal, or a spectral decomposition of A is available, then pCN satisfies all of our conditions for an efficient method.

Proposals for MALA and HMC are examples of proposals that are constructed by discretizing a stochastic differential equation that preserves the target distribution. Our theory allows us to consider a wider selection of possible AR(1) proposals for the MH algorithm, that are not necessarily based on discretizing a stochastic differential equation.

Appendix A

Proofs

A.1 Proof of Lemma 3.1.1

First note that

$$q(x, y) \propto \exp\left(\frac{1}{2}(My - Nx - \beta)^T(M + N)^{-1}(My - Nx - \beta)\right).$$

Simple algebra then yields

$$\begin{aligned} 2 \log \left(\frac{\pi_d(y)q(y, x)}{\pi_d(x)q(x, y)} \right) &= -y^T Ay + x^T Ax + 2b^T(y - x) \\ &\quad - (Mx - Ny - \beta)^T(M + N)^{-1}(Mx - Ny - \beta) \\ &\quad + (My - Nx - \beta)^T(M + N)^{-1}(My - Nx - \beta) \\ &= -y^T Ay + x^T Ax + 2b^T(y - x) \\ &\quad - ((M - N)(x + y))^T(M + N)^{-1}((M + N)(x - y)) \\ &\quad + 2\beta^T(M + N)^{-1}((M + N)(x - y)) \\ &= -y^T(A - \mathcal{A})y + x^T(A - \mathcal{A})x + 2(b - \beta)^T(y - x). \end{aligned}$$

A.1.1 Proof of Lemma 3.2.1

Suppose $\lim_{d \rightarrow \infty} (\sum_{i=1}^d |t_i|^r) / (\sum_{i=1}^d t_i^2)^{r/2} = 0$. Then for any $\epsilon > 0$ there exists a $D \in \mathbb{N}$ such that for any $d > D$, $\sum_{i=1}^d |t_i|^r < \epsilon (\sum_{i=1}^d t_i^2)^{r/2}$. Then for any $k \in \mathbb{N}$, taking $\epsilon = 2^{-r/2}$, there exists a $D \geq k$ such that for any $d > D$,

$$|t_k|^r \leq \sum_{i=1}^d |t_i|^r < \frac{1}{2^{r/2}} \left(\sum_{i=1}^d t_i^2 \right)^{r/2}.$$

Therefore, for any $d > D$, $t_k^2 < \frac{1}{2} \sum_{i=1}^d t_i^2$ and so

$$\frac{\sum_{i=1, i \neq k}^d |t_i|^r}{\left(\sum_{i=1, i \neq k}^d t_i^2 \right)^{r/2}} < \frac{\sum_{i=1}^d |t_i|^r}{\left(\frac{1}{2} \sum_{i=1}^d t_i^2 \right)^{r/2}} = 2^{r/2} \frac{\sum_{i=1}^d |t_i|^r}{\left(\sum_{i=1}^d t_i^2 \right)^{r/2}}.$$

A.2 Proof of Theorem 5.2.2

We use the following technical lemma in the proof of Theorem 5.2.2.

Lemma A.2.1 Suppose $\{t_i\} \subset \mathbb{R}$ is a sequence such that $0 < t_i \leq Cd^{-1/3}(\frac{i}{d})^{2\kappa}$ for $C > 0$ and $\kappa \geq 0$. If $s > 3$, then $\lim_{d \rightarrow \infty} \sum_{i=1}^d t_i^s = 0$.

Proof.

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d t_i^s \leq C^s \lim_{d \rightarrow \infty} d^{1-s/3} \sum_{i=1}^d \frac{1}{d} \left(\frac{i}{d}\right)^{2\kappa s} = C^s \lim_{d \rightarrow \infty} d^{1-s/3} \int_0^1 z^{2\kappa s} dz = 0.$$

■

Proof of Theorem 5.2.2. First note that VA and $V^{1/2}AV^{1/2}$ are similar, so they have the same eigenvalues. Lemma 5.2.1 implies that it is equivalent to study the MH algorithm with target and proposal given by (5.8). We now attempt to apply Theorems 4.1.4 and 4.2.1 to (5.8). Note that the splitting matrices for (5.8) are functions of $B = V^{1/2}AV^{1/2}$. Let $t_{d,i} = h\lambda_i = \mathcal{O}(d^{-t})$ for $t = 1/3 + \kappa$, let $s_i = h\lambda_i^2 = ld^{-1/3}(\frac{\lambda_i}{d})^{2\kappa} = \mathcal{O}(d^{-1/3})$ and let $\rho = (\theta - \frac{1}{2})/2$. Then

$$\begin{aligned} \tilde{\lambda}_i^2 &= (1 + \rho s_i)\lambda_i^2, & G_i &= 1 - \frac{\frac{1}{2}s_i}{1 + \frac{\theta}{2}s_i}, & \tilde{g}_i &= \frac{\frac{1}{2}s_i}{1 + \frac{\theta}{2}s_i}, & g_i &= \frac{s_i(1 + \rho s_i)}{1 + \frac{\theta}{2}s_i}, \\ r_i &= -\rho s_i, & \tilde{r}_i &= \frac{1}{1 + \rho s_i}, & \hat{r}_i &= 0, \end{aligned}$$

so that

$$\begin{aligned} T_{0i} &= T_{1i} = T_{2i} = 0, \\ T_{3i} &= \frac{-\frac{1}{2}\rho s_i^2(1 + \rho s_i)}{(1 + \frac{\theta}{2}s_i)^2}, & T_{4i} &= \frac{\frac{1}{2}\rho s_i^2}{(1 + \frac{\theta}{2}s_i)^2}, & T_{5i} &= \frac{\frac{1}{2}\rho s_i^{3/2}(1 - \frac{1-\theta}{2}s_i)}{(1 + \frac{\theta}{2}s_i)^2}. \end{aligned}$$

Hence

$$\tilde{g}_i^2 \hat{r}_i^2 \lambda_i = 0, \quad \tilde{g}_i^2 \lambda_i^{-1} = \mathcal{O}(t_{d,i} s_i) = o(t_{d,i}), \quad g_i \tilde{\lambda}_i^{-1} = \mathcal{O}(t_{d,i}), \quad \tilde{r}_i = \mathcal{O}(1)$$

as $d \rightarrow \infty$, and $T_{3i} = \mathcal{O}(d^{-2/3})$, $T_{4i} = \mathcal{O}(d^{-2/3})$ and $T_{5i} = \mathcal{O}(d^{-1/2})$ as $d \rightarrow \infty$.

We also have

$$\begin{aligned} \mu &= \lim_{d \rightarrow \infty} \sum_{i=1}^d \mu_i = \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{3i} + T_{4i} \\ &= \lim_{d \rightarrow \infty} -\frac{\rho^2}{2} \sum_{i=1}^d \frac{s_i^3}{(1 + \frac{\theta}{2}s_i)^2} = \lim_{d \rightarrow \infty} -\frac{\rho^2}{2} \sum_{i=1}^d s_i^3 = -\frac{l^6(\theta - \frac{1}{2})^2 \tau}{8}, \end{aligned}$$

and using Lemma A.2.1,

$$\begin{aligned}
\sigma^2 &= \lim_{d \rightarrow \infty} \sum_{i=1}^d \sigma_i^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{(1 + \frac{\theta}{2}s_i)^4} \left(\frac{1}{2}\rho^2 s_i^4 (1 + \rho s_i)^2 + \frac{1}{2}\rho^2 s_i^4 + \frac{1}{4}\rho^2 s_i^3 (1 - \frac{1-\theta}{2}s_i)^2 \right) \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d \left(\frac{1}{2}\rho^2 s_i^4 (1 + \rho s_i)^2 + \frac{1}{2}\rho^2 s_i^4 + \frac{1}{4}\rho^2 s_i^3 (1 - \frac{1-\theta}{2}s_i)^2 \right) \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{4}\rho^2 s_i^3 \\
&= \lim_{d \rightarrow \infty} \frac{l^6 (\theta - \frac{1}{2})^2 \tau}{4}.
\end{aligned}$$

Hence $\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + T_{2i}^2 + 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 < \infty$. Also note that, by Lemma A.2.1,

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{1i}^2 + T_{3i}^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{3i}^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{4}\rho^2 s_i^4 = 0.$$

It follows that $\frac{\mu}{\sigma} = -\sigma - \frac{\mu}{\sigma} = -l^3|\theta - \frac{1}{2}|\sqrt{\tau}/4$ and $\mu + \frac{\sigma^2}{2} = 0$. Hence we obtain (5.9) from Theorem 4.1.4, using Corollary 4.1.1.

For the expected jump size, first note that $\hat{r}_i = 0$. Also,

$$\frac{g_i}{\tilde{\lambda}_i^2} = \frac{h}{1 + \frac{\theta}{2}s_i} = h + \mathcal{O}(hs_i) = h + o(h), \quad \text{as } d \rightarrow \infty,$$

and using the fact that $\gamma_{d,i}^2$ is uniformly bounded (see proof of Theorem 4.1.4),

$$\frac{\tilde{g}_i^2 \gamma_{d,i}^{1/2}}{\lambda_i^2} = \mathcal{O}\left(\frac{s_i^2}{\lambda_i^2}\right) = \mathcal{O}(hs_i) = o(h) \quad \text{as } d \rightarrow \infty.$$

Therefore, applying Theorem 4.2.1 to the MH algorithm (5.8) we find

$$\mathbb{E}_{\pi_d}[(q_i^T(x'_i - x_i))^2] = 2h\Phi\left(-\frac{l^3|\theta - \frac{1}{2}|\sqrt{\tau}}{4}\right) + o(h)$$

as $d \rightarrow \infty$, where π_d is given in (5.8). Reversing the coordinate transformation $x \leftrightarrow V^{-1/2}x$ we obtain (5.10).

A.3 Proof of Lemma 5.3.1

By replacing π_d with π^* in the proof of Lemma 3.1.1 it follows that

$$\pi^*(x)q(x, y) = \pi^*(y)q(y, x)$$

where q corresponds the transition kernel for the SLA proposal. That is, $q(\cdot, \cdot)$ satisfies detailed balance with respect to $\pi^*(\cdot)$. It then follows from

$$q_L(x, y) = \int \cdots \int q(x, y^{(1)})q(y^{(1)}, y^{(2)}) \cdots q(y^{(L-1)}, y) dy^{(1)}dy^{(2)} \cdots dy^{(L-1)}$$

that $q_L(\cdot, \cdot)$ satisfies detailed balance with respect to $\pi^*(\cdot)$, so that

$$\pi^*(x)q_L(x, y) = \pi^*(y)q_L(y, x)$$

Using this fact, we find

$$\frac{\pi_d(y)q_L(y, x)}{\pi_d(x)q_L(x, y)} = \frac{\frac{\pi_d(y)}{\pi^*(y)}\pi^*(y)q_L(y, x)}{\frac{\pi_d(x)}{\pi^*(x)}\pi^*(x)q_L(x, y)} = \frac{\frac{\pi_d(y)}{\pi^*(y)}\pi^*(y)q_L(y, x)}{\frac{\pi_d(x)}{\pi^*(x)}\pi^*(y)q_L(y, x)} = \frac{\pi_d(y)\pi^*(x)}{\pi_d(x)\pi^*(y)}.$$

Hence, result.

A.4 Proof of Theorem 5.3.1

We prove this theorem by applying Theorems 4.1.4 and 4.2.1 to L -step SLA. First, let us check the conditions for these theorems. Let $G = I - \frac{h}{2}A$, $\Sigma = hI$ be the proposal matrices for the SLA proposal. These matrices are functions of A . It follows that G_L and Σ_L for the L -step SLA proposal are also functions of A and $G_L\Sigma_L$ is symmetric. It then follows from Corollary 2.1.1 that the splitting matrices for L -step SLA are also functions of A .

Let $t_{d,i} = h\lambda_i = \mathcal{O}(d^{-1/3-\kappa})$ and $s_i = h\lambda_i^2 = \mathcal{O}(d^{-1/3})$ as $d \rightarrow \infty$. Since \mathcal{A} and β for L -step SLA are the same as for SLA, from the proof of Theorem 5.2.2 we have $\tilde{\lambda}_i^2 = (1 - \frac{1}{4}s_i)\lambda_i^2$, $\hat{r}_i = 0$, $r_i = \frac{1}{4}s_i$ and $\tilde{r}_i = (1 - \frac{1}{4}s_i)^{-1} = \mathcal{O}(1)$.

Also, $G_i = 1 - \frac{1}{2}s_i$ so the eigenvalues of G_L are $G_i^L = (1 - \frac{1}{2}s_i)^L = 1 - \frac{L}{2}s_i + \mathcal{O}(s_i^2)$. Hence, $\tilde{g}_{Li} = 1 - G_i^L = \frac{L}{2}s_i + \mathcal{O}(s_i^2)$ and $g_{Li} = 1 - G_i^{2L} = Ls_i + \mathcal{O}(s_i^2)$.

It then follows that $\tilde{g}_i^2\hat{r}_i^2\lambda_i = 0$, $\tilde{g}_{Li}^2\lambda_i^{-1} = \mathcal{O}(s_i^2\lambda_i^{-1}) = \mathcal{O}(s_it_{d,i}) = o(t_{d,i})$, and $g_{Li}\tilde{\lambda}_i^{-1} = \mathcal{O}(t_{d,i})$.

We also have for L -step SLA,

$$T_{3i} = \frac{1}{8}Ls_i^2 + \mathcal{O}(s_i^3), \quad T_{4i} = -\frac{1}{8}Ls_i^2 + \mathcal{O}(s_i^3), \quad T_{5i} = -\frac{1}{4}s_i(Ls_i + \mathcal{O}(s_i^2))^{1/2}(1 + \mathcal{O}(s_i)),$$

so $T_{3i} = \mathcal{O}(d^{-2/3})$, $T_{4i} = \mathcal{O}(d^{-2/3})$ and $T_{5i} = \mathcal{O}(d^{-1/2})$.

We can also calculate, using Lemma A.2.1,

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{3i}^2 = \frac{1}{64}L^2 \lim_{d \rightarrow \infty} \sum_{i=1}^d s_i^4 = 0,$$

and similarly, $\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{4i}^2 = 0$, and as in the proof of Theorem 5.2.2 (calculating σ^2),

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d T_{5i}^2 = \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{16}Ls_i^3 = L\frac{l^6\tau}{16} =: \sigma_L^2,$$

Using Lemma A.2.1, and as for calculating μ in the proof of Theorem 5.2.2,

$$\begin{aligned}
\mu_L &= \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{3i} + T_{4i} = \lim_{d \rightarrow \infty} \sum_{i=1}^d -\frac{1}{2} r_i g_i (\tilde{r}_i - 1) \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d -\frac{1}{32} L s_i^3 + \mathcal{O}(s_i^4) \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d -\frac{1}{32} L s_i^3 \\
&= -L \frac{l^6 \tau}{32}.
\end{aligned}$$

Hence, we obtain (5.12) from Theorem 4.1.4 and Corollary 4.1.1.

For the expected jump size, (5.13) follows from $\hat{r}_i = 0$, $g_{Li} \tilde{\lambda}_i^{-2} = Lh + o(h)$, $\tilde{g}_{Li}^2 \lambda_i^{-2} = o(h)$, and $\gamma_{d,i}$ bounded uniformly (see proof of Theorem 4.1.4).

A.5 Proof of Theorem 5.4.1

The result will follow from Corollary 2.1.1 with $G = (K^L)_{11}$ and $\Sigma = (K^L)_{12} V^{-1} (K^L)_{12}^T$ but we must first check that $G\Sigma$ is symmetric. Define \mathcal{V} , \mathcal{K} and B as in Theorem 5.4.3. Then $K = \mathcal{V}\mathcal{K}\mathcal{V}^{-1}$, so that $K^L = \mathcal{V}\mathcal{K}^L\mathcal{V}^{-1}$ and

$$(K^L)_{11} = V^{1/2} (\mathcal{K}^L)_{11} V^{-1/2} \quad \text{and} \quad (K^L)_{12} = V^{1/2} (\mathcal{K}^L)_{12} V^{1/2}.$$

Then

$$G\Sigma = V^{1/2} (\mathcal{K}^L)_{11} (\mathcal{K}^L)_{12}^2 V^{1/2}$$

which is symmetric because V and B are symmetric and $(\mathcal{K}^L)_{11}$ and $(\mathcal{K}^L)_{12}$ are polynomials of B .

A.6 Proof of Corollary 5.4.1

First note that

$$(I - (K^L)_{11}) \mathcal{A}^{-1} \beta = \left(S J \begin{bmatrix} 0 \\ \frac{h}{2} b \end{bmatrix} \right)_1,$$

so we are required to show that

$$(I - (K^L)_{11}) \mathcal{A}^{-1} b = \left(S J \begin{bmatrix} 0 \\ \frac{h}{2} b \end{bmatrix} \right)_1,$$

which holds if

$$(I - K^L) \begin{bmatrix} A^{-1} b \\ 0 \end{bmatrix} = S J \begin{bmatrix} 0 \\ \frac{h}{2} b \end{bmatrix}.$$

Using $S = (I - K)^{-1} (I - K^L)$, we can equivalently show

$$(I - K) \begin{bmatrix} A^{-1} b \\ 0 \end{bmatrix} = J \begin{bmatrix} 0 \\ \frac{h}{2} b \end{bmatrix},$$

which is easy to check.

A.7 Proof of Theorem 5.4.2

Define a spectral decomposition

$$V^{1/2}AV^{1/2} = Q\Lambda Q^T \quad (\text{A.1})$$

where Q is an orthogonal matrix and $\Lambda = \text{diag}(\lambda_1^2, \dots, \lambda_d^2)$ is a diagonal matrix of eigenvalues of $V^{1/2}AV^{1/2}$ (VA is similar to $V^{1/2}AV^{1/2}$ so they have the same eigenvalues). Also define \mathcal{V} as in Theorem 5.4.3 and

$$\tilde{Q} = \begin{bmatrix} Q & 0 \\ 0 & Q \end{bmatrix} \in \mathbb{R}^{2d \times 2d}.$$

A similarity transform of K is defined by

$$K = \mathcal{V}\tilde{Q}\tilde{K}\tilde{Q}^T\mathcal{V}^{-1} \quad \text{with} \quad \tilde{K} = \begin{bmatrix} I - \frac{h^2}{2}\Lambda & hI \\ -h\Lambda + \frac{h^3}{4}\Lambda^2 & I - \frac{h^2}{2}\Lambda \end{bmatrix}.$$

Hence K and \tilde{K} have the same eigenvalues. Moreover, $K^L = \mathcal{V}\tilde{Q}\tilde{K}^L\tilde{Q}^T\mathcal{V}^{-1}$ and it follows that

$$(K^L)_{11} = V^{1/2}Q(\tilde{K}^L)_{11}Q^TV^{-1/2}.$$

Thus $(K^L)_{11}$ and $(\tilde{K}^L)_{11}$ are similar.

Notice that \tilde{K} is a 2×2 block matrix where each $d \times d$ block is diagonal. Therefore, \tilde{K}^L is also a 2×2 block matrix with diagonal blocks. In particular, $(\tilde{K}^L)_{11}$ is a diagonal matrix, so the eigenvalues of $(\tilde{K}^L)_{11}$ are on the diagonal of $(\tilde{K}^L)_{11}$. Moreover,

$$[(\tilde{K}^L)_{11}]_{ii} = (k_i^L)_{11}$$

where $[(\tilde{K}^L)_{11}]_{ii}$ is the i^{th} diagonal entry of $(\tilde{K}^L)_{11}$, $(k_i^L)_{11}$ is the $(1,1)$ entry of the matrix $k_i^L \in \mathbb{R}^{2 \times 2}$, and $k_i \in \mathbb{R}^{2 \times 2}$ is defined by

$$k_i = \begin{bmatrix} (\tilde{K}_{11})_{ii} & (\tilde{K}_{12})_{ii} \\ (\tilde{K}_{21})_{ii} & (\tilde{K}_{22})_{ii} \end{bmatrix} = \begin{bmatrix} 1 - \frac{h^2}{2}\lambda_i^2 & h \\ -h\lambda_i^2 + \frac{h^3}{4}\lambda_i^4 & 1 - \frac{h^2}{2}\lambda_i^2 \end{bmatrix}.$$

The matrix k_i can be factorized

$$k_i = \begin{bmatrix} 1 & 0 \\ 0 & a \end{bmatrix} \begin{bmatrix} \cos(\theta_i) & -\sin(\theta_i) \\ \sin(\theta_i) & \cos(\theta_i) \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & a^{-1} \end{bmatrix}$$

where $a = \lambda\sqrt{1 - \frac{h^2}{4}\lambda_i^2}$ and $\theta_i = -\cos(1 - \frac{h^2}{2}\lambda_i^2)$. Therefore,

$$k_i^L = \begin{bmatrix} 1 & 0 \\ 0 & a \end{bmatrix} \begin{bmatrix} \cos(L\theta_i) & -\sin(L\theta_i) \\ \sin(L\theta_i) & \cos(L\theta_i) \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & a^{-1} \end{bmatrix}$$

and hence

$$[(\tilde{K}^L)_{11}]_{ii} = (k_i^L)_{11} = \cos(L\theta_i).$$

A.8 Proof of Theorem 5.4.4

Theorem 5.4.3 implies that it is equivalent to study the MH algorithm with target and proposal given by (5.18), and since $(\mathcal{K}^L)_{11}$ and $(\mathcal{K}^L)_{12}$ are functions of B we can apply Theorems 3.1.2 and 3.2.1. Using the spectral decomposition (A.1) note that

$$(\mathcal{K}^L)_{ij} = Q(\tilde{K}^L)_{ij}Q^T \quad \text{for } i, j = 1, 2.$$

where \tilde{K} is defined in the proof of Theorem 5.4.2, and where it is shown that $(\tilde{K}^L)_{ij}$ is diagonal and

$$[(\tilde{K}^L)_{11}]_{ii} = \cos(L\theta_i)$$

where $\theta_i = -\cos^{-1}(1 - \frac{h^2}{2}\lambda_i^2)$. Similarly,

$$[(\tilde{K}^L)_{12}]_{ii} = -a_i^{-1} \sin(L\theta_i)$$

where $a_i = \lambda_i \sqrt{1 - \frac{h^2}{4}\lambda_i^2}$. Moreover, $\mathcal{A} = Q(\tilde{K}^L)_{12}^2(I - (\tilde{K}^L)_{11}^2)Q^T$ so that $\tilde{\lambda}_i^2 = -a_i^2$ and if we let $s_i = h^2\lambda_i^2 = l^2 d^{-1/2}(\frac{\lambda_i}{d^\kappa})^2 = \mathcal{O}(d^{-1/2})$, then

$$\begin{aligned} \tilde{\lambda}_i^2 &= \lambda_i^2(1 - \frac{1}{4}s_i), & G_i &= \cos(L\theta_i), & \tilde{g}_i &= 1 - \cos(L\theta_i), & g_i &= \sin^2(L\theta_i), \\ r_i &= \frac{1}{4}s_i, & \tilde{r}_i &= \frac{1}{1 - \frac{1}{4}s_i}, & \hat{r}_i &= 0. \end{aligned}$$

Note that we used Corollary 5.4.1 to show $\hat{r}_i = 0$. Then

$$\begin{aligned} T_{0i} &= T_{1i} = T_{2i} = 0, \\ T_{3i} &= \frac{1}{8}s_i \sin^2(L\theta_i), & T_{4i} &= -\frac{\frac{1}{8}s_i \sin^2(L\theta_i)}{1 - \frac{1}{4}s_i}, & T_{5i} &= -\frac{\frac{1}{8}s_i \sin(2L\theta_i)}{\sqrt{1 - \frac{1}{4}s_i}}. \end{aligned}$$

Using the trigonometric expansion $\cos^{-1}(1 - z) = \sqrt{2z} + \mathcal{O}(z^{3/2})$, and defining T' such that $L = \frac{T'}{h}$ we find

$$L\theta_i = L(-\sqrt{s_i} + \mathcal{O}(s_i^{3/2})) = -Ls_i^{1/2}(1 + \mathcal{O}(s_i)) = -T'\lambda_i(1 + \mathcal{O}(d^{-1/2}))$$

hence, there exists a function $T''(d)$ such that $L\theta_i = -T''\lambda_i$ and $T''(d) = T' + \mathcal{O}(d^{-1/2})$.

To apply Theorems 3.1.2 and 3.2.1 we need to check (3.4). For some $h > 0$, $cl^2(\frac{i}{d})^{2\kappa} \leq d^{1/2}s_i = l^2(\frac{\lambda_i}{d^\kappa})^2 \leq Cl^2(\frac{i}{d})^{2\kappa}$ and so we find

$$\begin{aligned} \lim_{d \rightarrow \infty} \frac{\sum_{i=1}^d |T_{3i}|^{2+\delta}}{\left(\sum_{i=1}^d |T_{3i}|^2\right)^{1+\delta/2}} &= \lim_{d \rightarrow \infty} \frac{\sum_{i=1}^d |s_i \sin^2(L\theta_i)|^{2+\delta}}{\left(\sum_{i=1}^d |s_i \sin^2(L\theta_i)|^2\right)^{1+\delta/2}} \\ &= \lim_{d \rightarrow \infty} d^{-\delta/2} \left(\frac{\left(\sum_{i=1}^d |d^{1/2}s_i \sin^2(T''\lambda_i)|^{2+\delta}\right)^{1/(2+\delta)}}{\left(\sum_{i=1}^d |d^{1/2}s_i \sin^2(T''\lambda_i)|^2\right)^{1/2}} \right)^{2+\delta} \\ &= 0 \end{aligned}$$

since $\|v\|_q \leq \|v\|_p$ for $q \geq p > 0$ for l^p -norms on \mathbb{R}^d . Similar arguments verify (3.4) for T_{4i} and T_{5i} . Now we can apply Theorem 3.1.2 with

$$\begin{aligned}
\mu &= \lim_{d \rightarrow \infty} \sum_{i=1}^d T_{3i} + T_{4i} \\
&= -\frac{1}{32} \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{s_i^2 \sin^2(L\theta_i)}{1 - \frac{1}{4}s_i} \\
&= -\frac{1}{32} \lim_{d \rightarrow \infty} \sum_{i=1}^d s_i^2 \sin^2(L\theta_i) \\
&= -\frac{l^4}{32} \lim_{d \rightarrow \infty} \frac{1}{d^{1+4\kappa}} \sum_{i=1}^d \lambda_i^4 \sin^2(\lambda_i T''(d)) \\
&= -\frac{l^4}{32} \lim_{d \rightarrow \infty} \frac{1}{d^{1+4\kappa}} \sum_{i=1}^d \lambda_i^4 \sin^2(\lambda_i T') \\
&= -\frac{l^4 \tau}{32}
\end{aligned}$$

and similarly,

$$\begin{aligned}
\sigma^2 &= \lim_{d \rightarrow \infty} \sum_{i=1}^d 2T_{3i}^2 + 2T_{4i}^2 + T_{5i}^2 \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{32} s_i^2 \sin^4(L\theta_i) + \frac{1}{32} \frac{s_i^2 \sin^4(L\theta_i)}{(1 - \frac{1}{4}s_i)^2} + \frac{1}{64} \frac{s_i^2 \sin^2(2L\theta_i)}{1 - \frac{1}{4}s_i} \\
&= \lim_{d \rightarrow \infty} \sum_{i=1}^d \frac{1}{16} s_i^2 \sin^4(L\theta_i) + \frac{1}{64} s_i^2 \sin^2(2L\theta_i) \\
&= \frac{1}{16} \lim_{d \rightarrow \infty} \sum_{i=1}^d s_i^2 \sin^2(L\theta_i) \\
&= \frac{l^4 \tau}{16}.
\end{aligned}$$

Hence $\frac{\mu}{\sigma} = -\sigma - \frac{\mu}{\sigma} = -\frac{l^2 \sqrt{\tau}}{8}$ and $\mu + \sigma^2/2 = 0$, so from Theorem 3.1.2 we obtain (5.19).

For the expected jump size, we apply Theorem 3.2.1 with

$$\begin{aligned}
U_1 &= \frac{\tilde{g}_i^2}{\tilde{\lambda}_i^2} + \frac{g_i}{\tilde{\lambda}_i^2} = \frac{(1 - \cos(L\theta_i))^2}{\lambda_i^2} + \frac{\sin^2(L\theta_i)}{\lambda_i^2(1 - \frac{1}{4}s_i)} = \frac{2(1 - \cos(\lambda_i T'))}{\lambda_i^2} + \mathcal{O}\left(\frac{s_i}{\lambda_i^2}\right) \\
&= \frac{2(1 - \cos(\lambda_i T'))}{\lambda_i^2} + \mathcal{O}(h)
\end{aligned}$$

as $d \rightarrow \infty$. Also, it is straightforward to show that $\mu_i = \mathcal{O}(d^{-1})$ and $\sigma_i^2 = \mathcal{O}(d^{-1})$, so μ^- and σ^- exist. Recall that $Z \sim \mathcal{N}(\mu, \sigma^2)$ and $X \sim \mathcal{N}(\mu^-, (\sigma^-)^2)$, then $X - Z \sim \mathcal{N}(-\mu_i, \sigma^2 + (\sigma^-)^2)$, and using the fact that $z \mapsto 1 \wedge e^z$ is monotone and globally Lipschitz continuous

$$|U_2 - \mathbb{E}[\alpha(x, y)]| = |\mathbb{E}[1 \wedge e^X] - \mathbb{E}[1 \wedge e^Z]| \leq |\mathbb{E}[X - Z]| = |\mu_i| = \mathcal{O}(d^{-1}).$$

Then $U_2 = \mathbb{E}[\alpha(x, y)] + \mathcal{O}(d^{-1})$ as $d \rightarrow \infty$, and

$$\begin{aligned}
|U_3| &= (\sigma_i^2 + \mu_i^2)^{1/2} \frac{\sqrt{3}}{\lambda_i^2} (\tilde{g}_i^2 + \tilde{r}_i g_i) \\
&= (\sigma_i^2 + \mu_i^2)^{1/2} \frac{\sqrt{3}}{\lambda_i^2} \left((1 - \cos(L\theta_i))^2 + \frac{1}{1 - \frac{1}{4}s_i} \sin^2(L\theta_i) \right) \\
&= \mathcal{O}(d^{-1/2}).
\end{aligned}$$


Therefore, we obtain (5.20).

References

- [1] V. I. Arnol'd. *Mathematical methods of classical mechanics*, volume 60 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1989. Translated from the Russian by K. Vogtmann and A. Weinstein.
- [2] Owe Axelsson. *Iterative solution methods*. Cambridge University Press, Cambridge, 1994.
- [3] Mylène Bédard. Weak convergence of metropolis algorithms for non-i.i.d. target distributions. *The Annals of Applied Probability*, 17(4):1222–1244, 2007.
- [4] Mylène Bédard. Optimal acceptance rates for metropolis algorithms: Moving beyond 0.234. *Stochastic Processes and their Applications*, 118(12):2198 – 2222, 2008.
- [5] Mylène Bédard and Jeffrey S. Rosenthal. Optimal scaling of Metropolis algorithms: heading toward general target distributions. *Canad. J. Statist.*, 36(4):483–503, 2008.
- [6] A. Beskos, N. Pillai, G. Roberts, J.M. Sanz-Serna, and A. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.
- [7] A. Beskos, F. J. Pinski, J. M. Sanz-Serna, and A. M. Stuart. Hybrid Monte Carlo on Hilbert spaces. *Stochastic Process. Appl.*, 121(10):2201–2230, 2011.
- [8] A. Beskos, G. Roberts, A. Stuart, and J. Voss. MCMC methods for diffusion bridges. *Stoch. Dyn.*, 8(3):319–350, 2008.
- [9] Alexandros Beskos, Gareth Roberts, and Andrew Stuart. Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions. *Ann. Appl. Probab.*, 19(3):863–898, 2009.
- [10] Alexandros Beskos and Andrew Stuart. Computational complexity of metropolis-hastings methods in high dimensions. In Pierre L' Ecuyer and Art B. Owen, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2008*, pages 61–71. Springer Berlin Heidelberg, 2009.
- [11] Alexandros Beskos and Andrew Stuart. MCMC methods for sampling function space. In *ICIAM 07—6th International Congress on Industrial and Applied Mathematics*, pages 337–364. Eur. Math. Soc., Zürich, 2009.
- [12] Patrick Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, third edition, 1995. A Wiley-Interscience Publication.
- [13] Sergio Blanes, Fernando Casas, and J. M. Sanz-Serna. Numerical integrators for the hybrid Monte Carlo method. *SIAM J. Sci. Comput.*, 36(4):A1556–A1580, 2014.
- [14] L.A. Breyer and G.O. Roberts. From metropolis to diffusions: Gibbs states and optimal scaling. *Stochastic Processes and their Applications*, 90(2):181 – 206, 2000.
- [15] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. MCMC methods for functions: modifying old algorithms to make them faster. *Statist. Sci.*, 28(3):424–446, 2013.

- [16] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.
- [17] Colin Fox. Polynomial accelerated mcmc and other sampling algorithms inspired by computational optimization. In Josef Dick, Frances Y. Kuo, Gareth W. Peters, and Ian H. Sloan, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2012*, volume 65 of *Springer Proceedings in Mathematics & Statistics*, pages 349–366. Springer Berlin Heidelberg, 2013.
- [18] Colin Fox and Albert Parker. Convergence in variance of Chebyshev accelerated Gibbs samplers. *SIAM J. Sci. Comput.*, 36(1):A124–A147, 2014.
- [19] Colin Fox and Albert Parker. Accelerated Gibbs sampling of normal distributions using matrix splittings and polynomials. *Bernoulli*; in the press, 2016.
- [20] Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- [21] Jonathan Goodman and Alan D. Sokal. Multigrid monte carlo method. conceptual foundations. *Phys. Rev. D*, 40:2035–2071, Sep 1989.
- [22] Peter J. Green and Xiao-liang Han. Metropolis methods, Gaussian proposals and anti-thetic variables. In *Stochastic models, statistical methods, and algorithms in image analysis (Rome, 1990)*, volume 74 of *Lecture Notes in Statist.*, pages 142–164. Springer, Berlin, 1992.
- [23] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 2006. Structure-preserving algorithms for ordinary differential equations.
- [24] P. Hall and C. C. Heyde. *Martingale limit theory and its application*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1980. Probability and Mathematical Statistics.
- [25] Jun S. Liu. *Monte Carlo strategies in scientific computing*. Springer Series in Statistics. Springer-Verlag, New York, 2001.
- [26] Radford M Neal. Probabilistic inference using markov chain monte carlo methods. Technical Report CRG-TR-93-1, Department of Computer Science, University of Toronto Toronto, Ontario, Canada, 1993.
- [27] Jorge Nocedal and Stephen J. Wright. *Numerical optimization*. Springer Series in Operations Research. Springer-Verlag, New York, 1999.
- [28] Natesh S. Pillai, Andrew M. Stuart, and Alexandre H. Thiéry. Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. *Ann. Appl. Probab.*, 22(6):2320–2356, 2012.
- [29] G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Probab.*, 7(1):110–120, 1997.
- [30] Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 60(1):255–268, 1998.
- [31] Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, 16(4):351–367, 2001.

- [32] Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- [33] Pranab K. Sen and Julio M. Singer. *Large sample methods in statistics: an introduction with applications*. Chapman & Hall, New York, 1993.
- [34] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 19:451–559, 2010.



Electronics Group
Department of Physics
University of Otago
elec.otago.ac.nz

