

A NEW METHOD FOR COUPLING RANDOM FIELDS

L. A. BREYER AND G. O. ROBERTS

Abstract

Given a Markov chain, a stochastic flow that simultaneously constructs sample paths started at each possible initial value can be constructed as a composition of random fields. Here, a method is described for coupling flows by modifying an arbitrary field (consistent with the Markov chain of interest) by an independence Metropolis–Hastings iteration. The resulting stochastic flow is shown to have many desirable coalescence properties, regardless of the form of the original flow.

1. Introduction

In this note, we describe a novel approach to the problem of coupling random fields and, by extension, Markov processes. The problem setting is as follows. Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, upon which is defined a family $(\mathfrak{X}(z) : z \in Z)$ of random variables (called a *random field*), Z being some index set and each variable $\mathfrak{X}(z)$ taking values in some common measurable space (E, \mathcal{E}) . Under measurability conditions, the marginal distributions of the field are fully described by a kernel $P : Z \times \mathcal{E} \rightarrow [0, 1]$, in the sense that

$$\mathbb{P}(\mathfrak{X}(z) \in dy) = P(z, dy), \quad z \in Z. \tag{1}$$

Conversely, if we start with a kernel P and take definition (1) as a hypothesis, we can construct some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a family $X(z) : \Omega \rightarrow E$ ($z \in Z$) such that (1) holds. Such a probability is often called a *coupling*, at least when Z consists of two points only. The choices for \mathbb{P} are limitless, and it therefore becomes interesting to ask whether a probability space and a family $\mathfrak{X} = (\mathfrak{X}(z) : z \in Z)$ can be constructed which additionally satisfy some previously agreed condition.

(a) *The case $|Z| = 2$.* When the index set Z consists of only two (distinct) points, the problem that we shall deal with can be phrased as follows: given two marginal distributions $P_1(dy)$ and $P_2(dy)$ on a common space E , find a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a realization (X_1, X_2) such that

$$\mathbb{P}(X_1 = X_2) > 0, \quad X_1 \sim P_1(\cdot), X_2 \sim P_2(\cdot). \tag{2}$$

This problem can be solved [6] if and only if there exists some nontrivial measure $\nu(\cdot)$ on E such that both $P_1 \geq \nu$ and $P_2 \geq \nu$, in which case there exists a *maximal* probability \mathbb{P} satisfying (2) on a canonical probability space Ω . This result has been the basis for a very successful method, originally proposed by Nummelin, for coupling Markov chains [7].

(b) *The case $|Z| > 2$.* When the indexing set Z consists of more than two distinct points, we must replace (2) by a more suitable requirement on \mathbb{P} . There is no universal definition extending (2), but the following one seems natural.

Received 15 December 1999, revised 21 December 2001; published 24 October 2002.

2000 Mathematics Subject Classification 60J (primary); 60F (secondary).

© 2002, L. A. Breyer and G. O. Roberts

DEFINITION 1. Let Z be some index set, and consider a kernel $P : Z \times \mathcal{E} \rightarrow [0, 1]$ of probability measures on some measurable space (E, \mathcal{E}) . We say that a family $(\mathfrak{X}(z) : z \in Z)$ of random variables with values in E , defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is a *finite coupling of the field determined by P* if both (1) holds and

$$\mathbb{P}(z \mapsto \mathfrak{X}(z) \text{ has a finite range in } E) = 1.$$

The family is a *discrete coupling of the field determined by P* if (1) holds and

$$\mathbb{P}(z \mapsto \mathfrak{X}(z) \text{ has a discrete range in } E) = 1.$$

In our context, of course, we shall be interested in the case where $Z = E$.

One aim will be to find conditions on the kernel P such that there exists a finitely (or, respectively, discretely) coupled field. Interest in finite couplings is due to the recent developments in the literature on perfect simulation, which was spawned by the seminal work of Propp and Wilson [9]. The proof of existence that we offer is constructive, and is susceptible to implementation on a computer.

The construction of Markov chains for simulation purposes is usually accomplished by the use of a stochastic recursive representation of the Markov chain transitions [1]. This in turn implies a succession of stochastic fields, mapping at each simulation epoch the current state z of the chain into its subsequent value $\mathfrak{X}(z)$. On composition of these fields, we obtain a stochastic flow describing the simultaneous dynamics of the chain from all possible initial states.

It is of interest (especially in the fast-growing area of perfect simulation) to construct this flow in such a way as to force the coalescence of chains started from all possible initial states. Unfortunately, many natural stochastic recursive constructions do not lead to this, and it would be desirable to exhibit a general recipe for the modification of an arbitrary stochastic flow (consistent with this particular Markov chain dynamic) in order to guarantee the desired coalescence property. This is the motivation behind the present paper.

Our approach will be to modify the individual component fields of an existing flow, and to recombine these new fields into a flow that has improved coalescence properties. We shall use a method of modification based on the independence sampler (as commonly used in MCMC). Our first two main results, Theorems 2 and 5, deal with the modified fields themselves, showing that under suitable uniform ergodicity assumptions, the fields, $\mathfrak{X}(z)$, can be constructed so as to have finite image with arbitrarily high probability (Theorem 2) and with probability 1 by a suitable perfect simulation construction (Theorem 5).

Theorems 6 and 8 deal with the whole flow as constructed from the modified fields, $\mathfrak{X}^{(1)}(z), \mathfrak{X}^{(2)}(z), \dots$, say. Theorem 6 ensures the coupling of chains started at any two arbitrary initial values. Theorem 8 extends this result to an arbitrary collection of starting values. Finally, in Proposition 9, we focus on the problem of constructing perfect simulation algorithms for Metropolis–Hastings chains. These algorithms are not covered by our previous results, although their specific structure allows specialised results to be proved. Examples include the random walk Metropolis algorithm and the Langevin algorithm.

To see our field coupling methodology in action, the reader is advised to visit the URL given in Appendix A to see an online illustration.

2. A Markov chain on field space

The construction of the finitely (or, respectively, discretely) coupled field will be accomplished by applying ideas originally devised by Propp and Wilson [9] to a

Metropolis–Hastings-type chain whose states are represented by candidate fields. The successively produced fields in this chain will have progressively ‘smaller’ ranges until such time as only finite range realizations of the field are produced. This will almost surely occur after a finite number of iterations, uniformly fast in the ‘seed’ field. By the technique of ‘coupling from the past’, we generate from this chain a finitely coupled field.

We proceed to set up some notation, and then we describe the transitions of the above-mentioned Markov chain. Let (E, \mathcal{E}, μ) be a σ -finite measure space, and let Z be a topological space. We shall denote by $\mathfrak{F}(Z; E)$ the space of Borel measurable functions of Z into E . This will be thought of as the state space for a Markov chain \mathfrak{X}_n , defined below.

We also suppose given a probability kernel $P(z, dy) = p(z, y)\mu(dy)$ of Z into E , absolutely continuous with respect to μ . The existence of the density $p(z, y)$ here is mainly a matter of convenience, and can be dispensed with at the cost of small changes in the theory below.

For the construction of our Markov chain(s), we shall need an appropriate probability space (Ω, \mathcal{F}) , which we now describe. Take $\Omega = \mathfrak{F}(Z; E)$, and let \mathcal{F} be the σ -algebra generated by the function evaluations $z \mapsto \omega(z), \omega \in \Omega$. Define also the set Π of probability measures on (Ω, \mathcal{F}) which prescribe the marginals of the canonical random variable:

$$\Pi = \left\{ \mathbb{P} \in \mathcal{P}(\Omega) : \mathbb{P}(\omega : \omega(z) \in A) = P(z, A) \text{ for all } z \in Z, A \in \mathcal{B}(E) \right\}. \quad (3)$$

Note that Π is nonempty, since we can always take the product measure

$$\mathbb{P} = \bigotimes_z P(z, \cdot) \in \Pi.$$

However, in the interests of readability, given a random variable $\mathfrak{X} : \Omega \rightarrow \mathfrak{F}(Z; E)$, we shall often write ‘ $\mathfrak{X} \in \Pi$ ’ instead of the more cumbersome ‘the law of \mathfrak{X} belongs to Π ’.

We are now ready to define our Markov chain. In general, we shall need a family $\Theta_z(x, dy)$ of probability kernels satisfying $\Theta_z(x, \cdot) \gg P(x, \cdot)$ for each $z \in Z$; these will represent ‘proposals’. We suppose here also the existence of a density, $\Theta_z(x, dy) = \theta_z(x, y)\mu(dy)$.

To define the Markov chain (\mathfrak{X}_n) , we take $\mathfrak{X}_0 \in \Pi$ (‘any random variable \mathfrak{X}_0 whose law belongs to Π ’), and then use induction. Given $\mathfrak{X}_n = (\mathfrak{X}_n(z) : z \in Z)$, let $\Phi_z \sim \theta_z(\mathfrak{X}_n(z), \cdot)\mu(\cdot)$, and put $\mathfrak{X}_{n+1} = (\mathfrak{X}_{n+1}(z) : z \in Z)$, where

$$\mathfrak{X}_{n+1}(z) = \begin{cases} \Phi_z, & \text{if } p(z, \Phi_z)\theta_z(\Phi_z, \mathfrak{X}_n(z)) > \xi \cdot p(z, \mathfrak{X}_n(z))\theta_z(\mathfrak{X}_n(z), \Phi_z); \\ \mathfrak{X}_n(z), & \text{otherwise.} \end{cases} \quad (4)$$

Here, $\xi \sim U[0, 1]$ independently of \mathfrak{X}_n .

If we fix $z \in Z$ in the above, it is immediately seen that the process $n \mapsto \mathfrak{X}_n(z)$ is a Metropolis–Hastings chain with stationary distribution $P(z, \cdot)$ and proposal kernel $\Theta_z(x, dy)$. Since, moreover, by assumption $\mathfrak{X}_0(z) \sim p(z, \cdot)\mu(\cdot)$, one checks immediately (or see [11]) that

$$\mathbb{P}(\mathfrak{X}_n(z) \in dy) = p(z, y)\mu(dy), \quad n > 0; \quad \text{that is, } \mathfrak{X}_n \in \Pi. \quad (5)$$

However, when viewed as a family indexed by z , the chains $\mathfrak{X}_n(\cdot)$ are not independent as z varies throughout Z , since we reuse the same random number ξ .

In the examples below, we describe some choices of the kernel density $\theta_z(x, y)$ which guarantee that, for sufficiently large (but finite) n , the realization $z \mapsto \mathfrak{X}_n(z)$ almost surely has finite range. None of the choices that we list is claimed to be ‘optimal’.

EXAMPLE 1. The simplest example of such a chain is arguably obtained if we take $\theta_z(x, y) \equiv \theta(y)$, a function of y alone. We call this the *simple field coupler*. For this choice, the formula (4) simplifies to

$$\mathfrak{X}_{n+1}(z) = \begin{cases} \Phi \sim \theta(\cdot)\mu(\cdot), & \text{if } p(z, \Phi)\theta(\mathfrak{X}_n(z)) > \xi \cdot p(z, \mathfrak{X}_n(z))\theta(\Phi); \\ \mathfrak{X}_n(z), & \text{otherwise.} \end{cases} \quad (6)$$

Note first that if $z \mapsto \mathfrak{X}_k(z)$ has finite range, then every subsequent function $z \mapsto \mathfrak{X}_{k+l}(z)$ will have finite range also. Thus the set of states in $\mathfrak{F}(Z; E)$ with finite range forms an absorbing set for any chain defined by (6). Let us now show that this chain \mathfrak{X}_n indeed evolves towards a state with finite range. We shall explain in the next section how to generate a finitely coupled field using a perfect simulation technique.

Recall that a subset $C \subset Z$ is called *1-small* [10] for the kernel $P(x, dy)$ on E if there exists a probability density ν and a constant $\epsilon > 0$ such that

$$\inf_{x \in C} p(x, y)\mu(dy) \geq \epsilon \nu(y)\mu(dy). \quad (7)$$

THEOREM 2. Suppose that Z can be entirely covered by a finite union of 1-small sets C_1, \dots, C_N , and suppose we can choose a probability density θ in such a way that

$$\epsilon_i \nu_i(\cdot) \leq \inf_{z \in C_i} p(z, \cdot) \leq \sup_{z \in C_i} p(z, \cdot) \leq \gamma_i \theta(\cdot), \quad i = 1, \dots, N, \quad (8)$$

for positive numbers $\epsilon_1, \dots, \epsilon_N, \gamma_1, \dots, \gamma_N$. If we define a Markov chain \mathfrak{X}_n on Π by the prescription (6), then

$$\mathbb{P}(\mathfrak{X}_n \text{ has finite range eventually}) = 1.$$

Proof. Clearly, from (6), if z is fixed the process, $n \mapsto \mathfrak{X}_n(z)$ is a Metropolis–Hastings Markov chain on E (an independence sampler, in fact), with stationary distribution $p(z, \cdot)\mu(\cdot)$. Therefore we certainly have (5), since $\mathfrak{X}_0(z)$ is already stationary. Now, for every $z \in Z$, the chain $n \rightarrow \mathfrak{X}_n(z)$ must almost certainly move eventually. This will happen the first time from (6) at time n if and only if

$$(p(z, \Phi)\theta(\mathfrak{X}_n(z))) > \xi \cdot (p(z, \mathfrak{X}_n(z))\theta(\Phi)). \quad (9)$$

By (7) and (8), if $z \in C_k$, the event (9) contains the event

$$\epsilon_k \nu_k(\Phi) > \gamma_k \xi \theta(\Phi), \quad (10)$$

which will occur eventually for some pair (Φ, ξ) since the pairs are independent, by the Borel–Cantelli lemma. When it occurs at time n_k say, all Markov chains $\mathfrak{X}_{n_k}(z)$ for $z \in C_k$ will accept the *same* jump to Φ , reducing the range of $z \mapsto \mathfrak{X}_{n_k}(z)$ on C_k to a single point. Note that after this common jump, the range over C_k is likely to break into more than one point, since the stationary distribution of $\mathfrak{X}(z)$ depends on z . However, this procedure will never create more than a finite number of break points at each iteration. This analysis is valid on each of the small sets C_1, \dots, C_N that cover Z , and another application of the Borel–Cantelli lemma now shows that each of these sets must experience a common jump within a finite time $\tau = \max\{n_1, \dots, n_N\}$, beyond which the realization $\mathfrak{X}_{\tau+n}$ must always have a finite range. \square

The simple field coupler does not usually generate finitely coupled random fields in a fixed, finite number of iterations, n say, if the ‘seed’ field \mathfrak{X}_0 is not already finitely coupled. To state an analogy with the behaviour of ergodic Markov chains, we make the following definition.

DEFINITION 3. A random field $\mathfrak{X} = (\mathfrak{X}(z) : z \in Z)$ is called *finitely coupled within $\epsilon > 0$* if

$$\mathbb{P}(z \mapsto \mathfrak{X}(z) \text{ has finite range}) \geq 1 - \epsilon.$$

With this definition, we can paraphrase Theorem 2 as follows. For each $\epsilon > 0$, there exists n such that \mathfrak{X}_n is finitely coupled within ϵ . It is *not* true that \mathfrak{X}_n is finitely coupled (within $\epsilon = 0$) for *any* n . Thus, given a realization \mathfrak{X}_0 of a seed field, we can never be certain that the realization \mathfrak{X}_n generated by the simple field coupler has finite range. This is important, as sampling the field at a random time may lead to a biased sample, as we shall see in Section 3. We shall construct a finitely coupled field \mathfrak{X} by a perfect simulation technique in the next section. For practical implementation of this technique, it is not even necessary to be able to generate an initial field \mathfrak{X}_0 , since its shape over the small set C_k is completely forgotten once the event (10) occurs.

The bounds (8) are not needed to implement the simple field coupler. If they are available, however, we can obtain an estimate of the speed of convergence to a finitely coupled field within $\epsilon > 0$. By the proof of Theorem 2, the realization \mathfrak{X}_n has finite range as soon as the event (10) has occurred for each $k = 1, \dots, N$. If we denote by T_k the first hitting time of the set $\{(y, x) : \epsilon_k \nu_k(y) > \gamma_k x \theta(y)\}$ by the IID series $(\Phi^{(n)}, \xi^{(n)})$, we have

$$\mathbb{P}(\mathfrak{X}_n \text{ has finite range}) \geq \mathbb{P}(\max(T_1, \dots, T_N) \leq n),$$

and this is independent of \mathfrak{X}_0 . This bound can be explicitly computed in particular cases.

If Z is only covered by a countable collection of 1-small sets, we can still run the algorithm (6), but the realization \mathfrak{X}_n will generally not have a discrete range in a finite time.

We end our description of the properties of the simple field coupler with some comments about the assumptions of Theorem 2. Let us suppose that $Z = E$, so that we may think of the kernel P as the transition function of a Markov chain with state space E .

It is a fairly easy exercise (left to the reader) to show that *coverage of E by a finite number of 1-small sets* and *irreducibility* imply that the whole set E is *petite*, and *aperiodicity* then gives the whole set E as *small*. This last statement implies that P is *uniformly ergodic*.

It is therefore of interest to know if the converse holds. Unfortunately, this case is not so straightforward. Indeed, if P is uniformly ergodic, it does not follow that E can be covered by a finite, even countable, collection of 1-small sets.

As a counterexample, consider the random scan Gibbs sampler with uniform target on $E = [0, 1] \times [0, 1]$. Here, the kernel P is given explicitly by

$$P(x_1, x_2; dy_1 \times dy_2) = \frac{1}{2}(\delta_{x_1}(dy_1)dy_2 + \delta_{x_2}(dy_2)dy_1).$$

Since for $(x_1, x_2) \neq (x'_1, x'_2) \in [0, 1] \times [0, 1]$, $P(x_1, x_2; \cdot)$ and $P(x'_1, x'_2; \cdot)$ are mutually singular measures unless $x_1 = x'_1$ or $x_2 = x'_2$, it is impossible to partition E into even a *countable* collection of 1-small sets. However, E is still *small* (and hence P is uniformly ergodic), for every part of the state space can be reached in two steps:

$$P^2(x_1, x_2; dy_1 \times dy_2) \geq \frac{1}{2}dy_1 \times dy_2 = \epsilon \nu(dy_1 \times dy_2).$$

In this counterexample, P did not have a density with respect to some σ -finite product measure on E . However, the random scan Gibbs sampler on a discrete state space, $E = \mathbb{Z}^2$ say, does have a joint density $p(x, y)$, and still there is only a countable covering by 1-small sets. Note that with a strictly positive continuous density and a locally compact state space E , there always is at least a *finite* covering by 1-small sets (see Corollary 4 below).

Intuitively, what is needed for a finite covering by 1-small sets is that the chain be able to do large jumps in the state space. We offer the following sufficient condition (which implies that the whole state space Z is 1-small).

COROLLARY 4. *Suppose that $P(x, dy)$ has a strictly positive continuous density with respect to some σ -finite reference probability μ on Z . If Z is compact, or if Z is only locally compact and the limit $p(\infty, y) := \lim_{z \rightarrow \infty} p(z, y) > 0$ exists and defines a probability density, then*

$$\mathbb{P}(\mathfrak{X}_n \text{ has finite range in finite time}) = 1.$$

Proof. If Z is locally compact, we add a point at infinity, making it compact. Thus we can always assume that Z is compact. Thus $M \equiv \sup_{x,y} p(x, y) < \infty$. Now by the strict positivity of the density $p(x, y)$ the sets $U_n = \{(z, y) : M^{-1}p(z, y) \in (1/n, 2/n)\}$ entirely cover the compact set $Z \times Z$. Hence there exists a finite subcovering U_{n_1}, \dots, U_{n_k} say, and the sets $C_j = \{z : (z, y) \in U_{n_j}\}$ for $j = 1, \dots, k$ are 1-small and cover Z :

$$\begin{aligned} P(z, dy) &= p(z, y)\mu(dy) \geq \frac{1}{n} 1_{U_{n_j}}(y)\mu(dy) \\ &= \epsilon_j \nu_j(dy) \text{ if } z \in C_j, \quad j = 1, \dots, k. \quad \square \end{aligned}$$

To illustrate the corollary, we suppose that P is the transition kernel associated with the random field on $Z = E = \mathbb{R}$ given by

$$\mathfrak{X}_0(z) = z + \sigma^2(z)\pi'(z)/2\pi(z) + \sigma(z)W, \quad W \sim \mathcal{N}(0, 1),$$

where we specify π and σ as follows. Let $\pi(x)$ be some function such that $\pi(x) \sim \exp(-\alpha|x|^\beta)$ as $|x| \rightarrow \infty$, where α and β are strictly positive real numbers. Taking any sequence $\gamma_s \geq 0$ such that $\gamma_s + \beta > 2$ for all s , let σ satisfy

$$\sigma^2(x)/(-\log \pi(x))^{\gamma_s/\beta} \rightarrow 1 \text{ as } |x| \rightarrow \infty.$$

It is shown in [12] that the kernel P associated with the field \mathfrak{X}_0 is uniformly ergodic, and it is left to the reader to check that the assumptions of the corollary above are verified. Thus Theorem 2 applies to the ‘seed’ field \mathfrak{X}_0 .

It is clear that the choice of proposal Φ for the simple field coupler will influence the acceptance rate of the algorithm. For any given $z_0 \in Z$, we can judiciously adapt the proposal density $\theta(\cdot)$ to the target density $p(z_0, \cdot)$, but this will be useless if the density $p(z, \cdot)$ varies highly as a function of z . This difficulty can be addressed very simply by implementing an auxiliary variable which can be used to tune the proposal θ . The next example shows one way to do this. It will be called the *independence field coupler*.

EXAMPLE 2. In this example, we present a version of the chain (4) which allows a range of proposal distributions, thereby better taking into account the shape of $p(z, \cdot)$. Referring back to (4), we suppose given a family $\theta_z(y)$ of proposal distributions. Typical choices of $\theta_z(\cdot)$ would be ‘close’ to $p(z, \cdot)$. We shall need a selection mechanism, to decide which distribution θ_{z_0} is to be used at every step of the algorithm. The simplest choice is to take z_0 distributed independently on Z according to some distribution $\eta(\cdot)$, perhaps assigning higher weight to highly variable or otherwise difficult regions of the field. We call the *independence field coupler* the following Markov chain (\mathfrak{X}_n) on $\mathfrak{F}(Z; E)$. Given \mathfrak{X}_n , draw independently $Y^{(n+1)} := Y \sim \eta(\cdot)$, and set

$$\mathfrak{X}_{n+1}(z) = \begin{cases} \Phi_Y \sim \theta_Y(\cdot)\mu(\cdot), & \text{if } p(z, \Phi_Y)\theta_Y(\mathfrak{X}_n(z)) > \xi \cdot p(z, \mathfrak{X}_n(z))\theta_Y(\Phi_Y); \\ \mathfrak{X}_n(z), & \text{otherwise.} \end{cases} \quad (11)$$

Note that we can think of the object $(Y^{(n)}, \mathfrak{X}_n(z) : z \in Z)$ as a new random field, where the index set is now $Z \cup \{z_0\}$ for some isolated point $z_0 \notin Z$, and $\mathfrak{X}_n(z_0) := Y^{(n)}$. With this notation, we can fit the rule (11) in the framework (4) (this is left as an exercise for the reader).

The auxiliary variables $Y^{(n)}$ are IID, and given $\mathfrak{X}_n \in \Pi$ (by the inductive hypothesis), we also have

$$\begin{aligned} \mathbb{P}(\mathfrak{X}_{n+1}(z) \in dx) &= \int \mathbb{P}(\mathfrak{X}_{n+1}(z) \in dx \mid Y = y)\eta(dy) \\ &= \int p(z, x)\mu(dx)\eta(dy) \quad (\text{using (5)}) \\ &= p(z, x)\mu(dx), \end{aligned}$$

since if $Y = y$ is given, we just have the previous simple field coupler (6) with $\theta(\cdot) = \theta_y(\cdot)$. Thus, clearly, $\mathfrak{X}_{n+1} \in \Pi$ if $\mathfrak{X}_n \in \Pi$. We can now repeat the arguments of the previous example, yielding identical theorems for this new field coupler if we simply replace (8) with the generalization

$$\epsilon_i \nu_i(\cdot) \leq p(z, \cdot) \leq \gamma_i \theta_y(\cdot), \quad z, y \in C_i, \text{ for } i = 1, \dots, N. \quad (12)$$

Let us assume this done.

The construction (11), while an improvement over (6), may suffer from another defect. Due to the choice of auxiliary variable Y independently of the current state of the field $\mathfrak{X}_n = (\mathfrak{X}_n(z) : z \in Z)$, we can expect some redundancy in the proposed values $\Phi^{(1)}, \Phi^{(2)}, \dots$ when $Y^{(n)}$ occurs in a location sufficiently close to a previous value $Y^{(k)}$ (where $k < n$), whose associated proposal $\Phi^{(k)}$ was accepted. In our simulations, this appeared to matter only when the choices of θ_z and η were ‘bad’. Clearly, the scope for further extensions of the field coupler (\mathfrak{X}_n) is limitless.

3. Generating finitely coupled fields

The field couplers of the previous section are straightforward to implement, and computer experiments suggest that a small number of iterations will suffice to obtain a finite range. Here we mean, of course, that the Markov chain \mathfrak{X}_n must be simulated for a small but fixed (non-random) number of iterations, starting from an uncoupled version of the field. However, for any n , the realized field \mathfrak{X}_n will only ever be finitely coupled within some $\epsilon > 0$.

For application to simulation problems, it is desirable to implement a procedure that will output *with certainty* a finitely coupled version of the field of interest. Here, we describe how this can be accomplished in general, while in the next section, some concrete Markov chain Monte Carlo examples will be discussed.

The difficulty encountered when trying to output a finitely coupled field using the Markov chain \mathfrak{X}_n is that, even though $\mathfrak{X}_0 \in \Pi$ and $\mathfrak{X}_n \in \Pi$ for all deterministic times $n \geq 0$ (recall that Π was defined in (3)), if we stop the chain at some random time τ by observing the field $(\mathfrak{X}_n(z) : z \in Z)$ and deciding that it has fully coupled, the result \mathfrak{X}_τ no longer belongs to Π in general. To illustrate, consider the following naive procedure.

Faulty reasoning: Suppose that $Z = [0, 1]$, with $p(z, y) = (z + 1)y^z$, $0 \leq y \leq x$. We consider the simple field coupler with proposal given by $\Phi \sim p(0, \cdot)$. Let τ be the first time that the field is finitely coupled. It is easy to check that at each iteration of the field coupler, the set of z values that accept the proposed move is of the form $[0, z_n^*]$ for

some random $0 < z_n^* \leq 1$. Given a suitable $U(0, 1)$ random variable W , we can set $\mathfrak{X}_0(z) = W^{(1+z)^{-1}}$, so that finite coupling will occur precisely when $z_n^* = 1$. Therefore $\mathfrak{X}_\tau(\cdot)$ is a constant map. However this is inconsistent with the marginal mean transition values: $\int_0^1 yp(x, y) = (x + 1)/(x + 2)$, so that \mathfrak{X}_τ does not belong to Π .

In general, stopping times τ for which $\mathfrak{X}_\tau(z) \sim P(z, \cdot)$ are hard to find. If both $\mathfrak{X}_\tau \in \Pi$ and \mathfrak{X}_τ is independent of τ , then τ is known as a *strong stationary time* (see [2]).

There is another method to generate a finitely coupled field in finite time, by going backwards in time rather than forward. The details will occupy the rest of this section. Throughout, we hypothesize that Z is covered by a finite number of 1-small sets C_1, \dots, C_N , and that (12) holds. We now briefly remind the reader of the Propp and Wilson formalism for *perfect simulation* by Markov chains, duly adapted to our requirements.

The equations (6) and (11) can be viewed as implicitly defining a sequence of independent random functions $G_n : \mathfrak{F}(Z; E) \times \Omega \rightarrow \mathfrak{F}(Z; E)$ so that $\mathfrak{X}_{n+1}(\omega) = G_n(\mathfrak{X}_n(\omega), \omega)$. If function composition occurs in the first variable, then we have equivalently

$$\mathfrak{X}_n(\omega) = (G_n \circ G_{n-1} \circ \dots \circ G_1)(\mathfrak{X}_0(\omega), \omega).$$

It is convenient to extend the family (G_n) for values $n = 0, -1, -2, -3, \dots$ so that we may think of the chain \mathfrak{X}_n as evolving from the distant past.

By (5), it is clear that for any $u \in \mathfrak{F}(Z; E)$,

$$(G_n \circ G_{n-1} \circ \dots \circ G_1)(u) \in \Pi, \quad \text{if } u \in \Pi. \tag{13}$$

But the maps G_n are exchangeable, so if we write $\mathfrak{Y}_n(u) = G_0 \circ G_{-1} \circ \dots \circ G_{-n}(u)$ for $u \in \Pi$, then by (13) we automatically have

$$\lim_{n \rightarrow \infty} \mathfrak{Y}_n(u) \in \Pi, \quad \text{if } u \in \Pi \text{ and the limit exists.} \tag{14}$$

Suppose, moreover, that we can prove (as we do in Theorem 5) that the random time

$$T = \inf\{n : \mathfrak{Y}_n(\cdot) \text{ is constant on } \Pi\}$$

is almost surely finite. In that case, we have

$$\mathfrak{Y}_n(\cdot) = \mathfrak{Y}_T \circ G_{T-1} \circ \dots \circ G_{-n}(\cdot) = \mathfrak{Y}_T(\cdot) \quad \text{on } \{T \leq n\}. \tag{15}$$

This fact, first noted by Propp and Wilson in a different context [9], forms the basis for their perfect simulation method. In our setting, it suffices now to construct a random time $T' \geq T$ for which we can conveniently test whether $\{T' \leq n\}$ has occurred. By (14) and (15) we must then have $\mathfrak{X}_0 \equiv \mathfrak{Y}_n(\cdot) \in \Pi$, a perfect specimen of a finitely coupled field. We shall do this now for the algorithm (11), as (6) can be viewed as the special case when $\theta_y(\cdot) \equiv \theta(\cdot)$ for all $y \in Z$.

Intuitively, we shall think of the field \mathfrak{X}_n defined by (11) as a random point pattern on $Z \times E$ (the graph of the map $z \mapsto \mathfrak{X}_n(z)$). One iteration of the independence field coupler (11) on the field \mathfrak{X}_n consists in the deletion of a section of the point pattern, replacing it by a new pattern. If we consider two initially distinct point patterns u and u' , then for any ℓ , the patterns $G_\ell(u)$ and $G_\ell(u')$ will agree at all points that were newly added, provided that these had been accepted by both patterns. Thus we can expect that for all sufficiently large n , all the points originally making up the patterns u and u' have disappeared from the patterns $G_0 \circ G_{-1} \circ \dots \circ G_{-n}(u)$ and $G_0 \circ G_{-1} \circ \dots \circ G_{-n}(u')$, which will therefore agree: $\mathfrak{Y}_n(u) = \mathfrak{Y}_n(u')$.

(FC-1):

Generate IID sequences $(\xi^{(0)}, Y^{(0)}, \Phi^{(0)})$, $(\xi^{(1)}, Y^{(1)}, \Phi^{(1)})$, \dots , where

$$\xi^{(k)} \sim U[0, 1], \quad Y^{(k)} \sim \eta(\cdot), \quad \Phi^{(k)} \sim \theta_{Y^{(k)}}(\cdot)\mu(\cdot), \quad k = 0, 1, 2, 3, \dots \quad (16)$$

(FC-2):

Let $T'_k = \min\{s \geq 0 : \epsilon_k \nu_k(\Phi^{(s)})/\theta_{Y^{(s)}}(\Phi^{(s)}) > \gamma_k \xi^{(s)}\}$

and set $T' = \max\{T'_k : k = 1, \dots, N\}$.

(FC-3):

For fixed n , define the random map G_{-n} from $(\xi^{(n)}, Y^{(n)}, \Phi^{(n)})$ by: if $u = (u(z) : z \in Z)$ is a field,

$$(G_{-n}(u))(z) = \begin{cases} \Phi^{(n)}, & \text{if } p(z, \Phi^{(n)})\theta_{Y^{(n)}}(u(z)) > \xi^{(n)} p(z, u(z))\theta_{Y^{(n)}}(\Phi^{(n)}); \\ u(z), & \text{otherwise.} \end{cases}$$

(FC-4):

Thus $\mathfrak{X} = G_0 \circ G_{-1} \circ \dots \circ G_{-T'}(u)$ is finitely coupled and belongs to Π . We output \mathfrak{X} .

Algorithm 1: The field coupler.

We now construct the stopping time $T' \geq T$. We assume that (12) holds. To generate a finitely coupled field, Algorithm 1 is used.

For step (FC-3), note that $(G_{-T'_k}(u))(z)$ is independent of u for $z \in C_k$, and hence so is $(\mathfrak{Y}_n(u))(z)$ for all $n \geq T'_k$. Consequently, $\mathfrak{Y}_{T'}(u)$ is a random field which is independent of u . Moreover, $\mathbb{P}(T' < \infty) = 1$ as follows from Theorem 5 below.

To summarize, we have the following result.

THEOREM 5. *Let Z be coverable by a finite number of 1-small sets C_1, \dots, C_N such that (12) holds. There exists a finitely coupled field \mathfrak{X} corresponding to the kernel P which almost surely can be constructed in a finite number of iterations of Algorithm 1.*

Proof. All we need to check is that $\mathbb{P}(T' < \infty) = 1$. Now each T'_k defined in (FC-2) is binomially distributed with strictly positive success probability, from which the claim follows. □

At each iteration, we can test very simply whether T' has occurred or not. The number $T' - T$ of iterations ‘wasted’ will depend upon the overall quality of the estimates (12). We shall describe some of our findings in the next section.

Note also that the total number of values in the range of \mathfrak{X} is bounded by T (and hence by T' , which has geometric tails).

Algorithm 1 given above is not *interruptible*. Bias is thus possible when a strategy of abandoning particularly long unsuccessful runs is adopted (as is invariably necessary in practice). Thus, for instance, if we stop the run when a fixed number n_0 of iterations has been reached, and begin afresh, we are biasing the output \mathfrak{X} in favour of fields that require at most n_0 iterations to generate. Since n_0 is also the maximum number of distinct values contained in the range of $z \mapsto \mathfrak{X}(z)$, we are effectively biasing in favour of a bounded range. While this is a potential problem, the stopping time T' will often be very small unless the estimates (12) are poor.

Fill [3] proposed another method for the perfect simulation of a target probability density, based on rejection sampling. This has the advantage of *being* interruptible. However, the chain $(\mathfrak{X}_n, Y^{(n)})$ is not reversible on Π (since the set of finitely coupled fields is absorbing) and a modification of Fill’s method for field coupling appears more complicated.

4. Flows and Markov chains: coupling two chains

For the remainder of this paper, we shall take $Z = E$. We are given a kernel $P(x, dy)$ on E representing the transition function of some Markov chain. For any sequence of random fields $\mathfrak{X}^{(1)}, \mathfrak{X}^{(2)}, \dots \in \Pi$, we can define a corresponding flow $F_{s,t} : E \rightarrow E$,

$$F_{s,t}(x) = \mathfrak{X}^{(t-1)} \circ \dots \circ \mathfrak{X}^{(s)}(x), \quad s < t, x \in E. \tag{17}$$

Then if we set $X_t(x) = F_{0,t}(x)$, we have simultaneously realized all Markov chains on E with transition function P and deterministic initial condition. It is common in simulation studies to take the random fields in the form $\mathfrak{X}^{(k)}(z) = f(z, \zeta_k)$, where f is some measurable function and (ζ_k) is an IID sequence. This has been called a stochastic recursive sequence (SRS) construction [1]. It should be noted, however, that it is by no means necessary to take the *same* function f for each k , so long as the resulting field $\mathfrak{X}^{(k)}$ always belongs to Π .

We shall now briefly discuss an application of the field couplers to the problem of coupling two such Markov chains, $X_t(x_1)$ and $X_t(x_2)$ say, both with the same transition kernel P .

Recall first that a *coupling time* here is a random time $T \geq 0$ such that $X_T(x_1) = X_T(x_2)$. Since both chains are built from the same flow, this also means that $X_{T+t}(x_1) = X_{T+t}(x_2)$ for all $t > 0$. The coupling time is called *successful* if $\mathbb{P}(T < \infty) = 1$. The existence of a successful coupling time is by no means guaranteed, and obviously depends on the choice of flow $F_{s,t}$.

Remark. Some flows do not have successful coupling times, even for uniformly ergodic kernels P . As an example, consider the random walk on the unit circle $S^1 = \{e^{i\theta} : \theta \in \mathbb{R}\}$, given by the flow

$$F_{t,t+1}(\exp(ix)) = \exp i(x + W_t), \quad W_t \sim \mathcal{N}(0, 1).$$

The Markov chain $X_t(e^{ix})$ is aperiodic and uniformly ergodic, with stationary distribution given by the Lebesgue measure on S^1 . From the description of $F_{t,t+1}$ it is also clear that two chains $X_t(\exp(ix_1))$ and $X_t(\exp(ix_2))$ evolve in parallel, and can never couple unless $x_1 = x_2 \pmod{2\pi}$.

The above example does not preclude the existence of some alternative flow $\tilde{F}_{t,t+1}$ which will allow a successful coupling to take place (indeed, the uniform ergodicity guarantees this existence). Now given *any* flow of the type in (17), we claim that the field couplers can generate such a new flow $\tilde{F}_{s,t}$, for which the coupling of the associated chains $\tilde{X}_t(x_1)$ and $\tilde{X}_t(x_2)$ is successful.

Indeed, suppose that for each s , we iterate the simple field coupler n times with seed $\mathfrak{X}^{(s)}$, obtaining a new field $\mathfrak{X}_n^{(s)} \in \Pi$. We define the flow

$$\tilde{F}_{s,t}(z) = \mathfrak{X}_n^{(t-1)} \circ \dots \circ \mathfrak{X}_n^{(s)}(z), \tag{18}$$

and corresponding Markov chains $\tilde{X}_t(z) = \tilde{F}_{0,t}(z)$. We shall consider only fields $\mathfrak{X}_n^{(s)}$ generated by the simple field coupler here, for simplicity and because the independence field coupler can be treated in a very similar way.

The following result shows that if the proposal density θ has heavy enough tails, then coupling of $\tilde{X}_t(x_1)$ and $\tilde{X}_t(x_2)$ is always successful for uniformly ergodic kernels P . This will be used in the next section.

Note that the coupling method proposed here does not require any explicit computations of minorization conditions, but merely the knowledge that they exist. This is unlike the classical coupling method based on splitting, whereby two chains both have to enter an explicitly defined small set C simultaneously, right before coupling.

The proof below also shows that if P is only positive Harris recurrent, then coupling is successful with strictly positive probability.

THEOREM 6. *Let $F_{s,t}$ be an arbitrary flow for P , and suppose that the target density $p(x, y)$ and proposal density $\theta(y)$ satisfy $p(z, y) < \gamma(z)\theta(y)$ for some function $\gamma(z)$ which is bounded on a 1-small set C . If P is aperiodic and uniformly ergodic, then*

$$\mathbb{P}(\tilde{X}_t(x_1) = \tilde{X}_t(x_2) \text{ eventually}) = 1, \quad \text{for all } x_1, x_2 \in E.$$

Proof. Since $F_{s,t}$ is arbitrary, we shall assume that (18) holds with $n = 1$. Recall that we denote by $X_s(x)$ the Markov chain started from x and constructed from the original flow $F_{s,t}$, and by $\tilde{X}_s(x)$ the Markov chain started from x constructed from the new flow $\tilde{F}_{s,t}$ obtained after iterating the simple field coupler. We shall analyse the law of

$$\tilde{T}(x_1, x_2) = \min \{s \geq 0 : \tilde{X}_s(x_1) = \tilde{X}_s(x_2)\}.$$

Since P is uniformly ergodic, there exist m and $\alpha > 0$ such that $\mathbb{P}(X_m(x) \in C) \geq \alpha$ holds for all $x \in E$. Denote by $H(y)$ the first exit time of the chain $(\Phi^{(s-1)}, \Phi^{(s)}, \xi^{(s)})$, started initially at $(y, \Phi^{(1)}, \xi^{(1)})$, from the set $\{(u, v, w) : p(u, v) > w\gamma(u)\theta(v)\}$. If $H(y) > m$, then because the proposal $\theta(\cdot)$ dominates $p(z, \cdot)$, we must have

$$\begin{aligned} p(y, \Phi^{(1)}) &> \gamma(y)\xi^{(1)}\theta(\Phi^{(1)}), \\ p(\Phi^{(1)}, \Phi^{(2)}) &> \gamma(\Phi^{(1)})\xi^{(2)}\theta(\Phi^{(2)}), \\ &\vdots \\ p(\Phi^{(m-1)}, \Phi^{(m)}) &> \gamma(\Phi^{(m-1)})\xi^{(m)}\theta(\Phi^{(m)}), \end{aligned}$$

and consequently $\tilde{X}_1(y) = \Phi^{(1)}, \tilde{X}_2(y) = \Phi^{(2)}, \dots, \tilde{X}_m(y) = \Phi^{(m)}$. In particular, either $\tilde{X}_m(x) = \tilde{X}_m(y) = \Phi^{(m)}$ (when $\tilde{T}(x, y) \leq m$ holds), or else $\tilde{X}_s(x) \neq \Phi^{(s)}$ for all $s \leq m$ (when $\tilde{T}(x, y) > m$ holds), and then $\tilde{X}_s(x) = X_s(x)$.

Now consider the event $B_{m+1} = \{\epsilon v(\Phi^{(m+1)}) > \max_{z \in C} \gamma(z)\xi^{(m+1)}\theta(\Phi^{(m+1)})\}$. When this holds, because the proposal $\theta(\cdot)$ dominates $p(z, \cdot)$, all chains that satisfy $\tilde{X}_m(x) \in C$ must couple; that is, $\tilde{X}_{m+1}(x) = \Phi^{(m+1)}$. Consequently,

$$\begin{aligned} \mathbb{P}(\tilde{T}(x, y) \leq m + 1) &\geq \mathbb{P}(X_m(x) \in C, \tilde{X}_m(y) \in C, H(y) > m, B_{m+1}) \\ &= \mathbb{P}(X_m(x) \in C)\mathbb{P}(\Phi^{(m)} \in C, H(y) > m)\mathbb{P}(B_{m+1}). \end{aligned}$$

Let A be a set of strictly positive stationary measure, and let $\beta > 0$ be a number such that $\inf_{y \in A} \mathbb{P}(\Phi^{(m)} \in C, H(y) > m) > \beta$ holds. Such a set must exist, for otherwise we would have $\tilde{F}_{s,s+1} \equiv F_{s,s+1}$ almost surely. We estimate

$$\inf_{y \in A} \mathbb{P}(\tilde{T}(x, y) \leq m + 1) \geq \alpha\beta\mathbb{P}(B_{m+1}) = \delta > 0,$$

and note that this is independent of x . To finish the proof, observe that if we begin with two chains $\tilde{X}_s(x_1)$ and $\tilde{X}_s(x_2)$ such that $x_1 \neq x_2$, then, because $\tilde{X}_s(x_2)$ is positive recurrent, it

must enter the set A infinitely often. Each time this occurs, regardless of the current location x of the chain $\tilde{X}_s(x_1)$, there exists at least δ probability that the two chains will couple within $m + 1$ steps. The Borel–Cantelli lemma therefore gives $\mathbb{P}(\tilde{T}(x_1, x_2) < \infty) = 1$. \square

In the case where the kernel density $p(x, y)$ and proposal $\theta(y)$ are continuous, on a locally compact state space E , a sufficient condition for the existence of $\gamma(z)$ is that $\lim_{y \rightarrow \infty} p(z, y)/\theta(y) = 0$ for all z .

Theorem 6 is also valid for the independence field coupler (defined in Section 2, Example 2). In this case, since there is a range of proposal densities $\theta_z(\cdot)$ to choose from, the hypothesis can be weakened as follows. There exists some function $\gamma(z, w)$ which is bounded on a set $C \times D$ such that $p(z, y) < \gamma(z, w)\theta_w(y)$, with C small and D of positive η -measure.

5. Coupling many Markov chains

In this section, we are interested in the question of whether the whole collection $X_t(x) : x \in E$ of Markov chains can be coupled in a common finite time. We already know from Theorem 6 that any two given chains $X_t(x_1)$ and $X_t(x_2)$ will couple successfully in a time $T(x_1, x_2) < \infty$ almost surely, but the possibility remains that $\sup_{x_1, x_2} T(x_1, x_2) = \infty$. Indeed, it was shown by Foss and Tweedie [4] that there exists *some* flow $F_{s,t}$ corresponding to P such that $\sup_{x_1, x_2} T(x_1, x_2) < \infty$ if and only if P is *uniformly ergodic*. Thus we must restrict ourselves here too to such kernels P . We would like to investigate the case that $F_{s,t}$ is fixed to be of the form (18) with $n \geq 1$.

DEFINITION 7. Let $F_{s,t}$ be a flow for the transition kernel P on the state space E . We say that $F_{s,t}$ *collapses in finite time* if

$$\mathbb{P}(x \mapsto F_{s,t}(x) \text{ is constant for all large values of } t) = 1.$$

We say that $F_{s,t}$ *thins at time* $r > s$ if

$$|\{F_{s,r+1}(x) : x \in E\}| < |\{F_{s,r}(x) : x \in E\}|.$$

As a direct consequence of the previous sections, we can state the next theorem.

THEOREM 8. Let P be uniformly ergodic with a density $p(x, y)$ which satisfies (8) for some collection C_1, \dots, C_N of 1-small sets covering E . If $F_{s,t}$ is any random flow with one-step transition probability P , then the flow $\tilde{F}_{0,t}$ defined by (18) using the simple field coupler with $n \geq N$ (where $\tilde{\mathfrak{X}}^{(t)} = F_{t,t+1}$) collapses in finite time.

Proof. Applying the simple field coupler $n \geq N$ times will ensure that the random field $\mathfrak{X}_n^{(s)}$ is finitely coupled within some $\epsilon > 0$. Consequently, the Borel–Cantelli lemma implies that the field \tilde{F}_{0,t_0} has finite range almost surely, for some finite t_0 . Now using Theorem 6 ensures that the flow \tilde{F}_{0,t_0+t} thins repeatedly until collapse. \square

Variations on the above theme are possible. For example, we can generate (perfectly) finitely coupled fields $\mathfrak{X}^{(s)}$ and use these to construct $\tilde{F}_{s,t}$. This is discussed below for Metropolis–Hastings chains.

In principle, we can now collapse the flows of any uniformly ergodic Markov chains, provided that the transition kernel P has a density. We give two examples.

EXAMPLE 3 (RANDOM WALK). Let $P(z, dy) = q(y - z)dy$, where $q(\cdot)$ is a continuous, not necessarily symmetric density on \mathbb{R}^d . A random field belonging to Π is easily found, namely

$$\mathfrak{X}(z) = z + W, \quad \text{where } W \sim q(y) dy.$$

We propose to find a finitely coupled version of $\tilde{\mathfrak{X}}$ on a compact subset $Z \subset \mathbb{R}^d$. We cover Z by a finite number of closed balls $C_i = B(z_i, r_i)$ say, $i = 1, \dots, N$, and such that $\epsilon_i v_i(y) = \inf_{z \in C_i} q(y - z)$ is not identically zero. Let g be the (continuous) density of a distribution on \mathbb{R}^d with heavier tails than q . We shall need to assume that $q(\cdot) \leq Cg(\cdot)$ for some constant C . Then we put

$$\theta_z(y) \equiv \theta_i(y) := g(y - z_i), \quad \text{if } z \in C_i = B(z_i, r_i),$$

and calculate the constants $\gamma_i = \max_{z \in C_i, y \in \mathbb{R}^d} q(y - z)/g(y - z_i)$. Finally, we take $Y \sim \eta(\cdot)$, where η is the uniform distribution on the set of points $\{z_1, \dots, z_N\}$. Now applying Algorithm 1, we obtain a new field $\tilde{\mathfrak{X}}(z) \sim q(y - z)dy$ which is almost surely finitely coupled on Z .

Note that we have not assumed that $q(\cdot)$ is symmetric. For symmetric and unimodal increment distributions, an alternative field coupling method for \mathfrak{X} exists, proposed by Murdoch and Green [8]. Their method, called the *bisection coupler*, uses translations and reflections to construct a finitely coupled field, and appears to require fewer computations in general. However, while the number of points in the range of the resulting field is almost surely finite, they show that it has infinite expectation. In contrast, for the field coupler presented here, the range of the field $\tilde{\mathfrak{X}}$ never contains more points than the number of iterations T' required to generate it, which as a random number has geometric tails.

EXAMPLE 4 (TIME-DISCRETIZED DIFFUSION). On the interval $Z = [0, 1]$, consider the Langevin random field

$$\mathfrak{X}(z) = z + \sigma W + \sigma^2 \pi'(z)/2\pi(z), \tag{19}$$

where $\pi(z)$ is a strictly positive differentiable function on Z and W is a standard Gaussian random variable. Let us write

$$\tilde{p}(z, y) = \exp(-|y - z - \sigma^2 \pi'(z)/2\pi(z)|^2 / 2\sigma^2) \tag{20}$$

for the unnormalized density of $\mathfrak{X}(z)$, so that $\tilde{p}(z, y) = cp(z, y)$ for some fixed constant c . None of our calculations requires the normalization constant, c . For $i = 1, \dots, N$, we let $C_i = [(i - 1)/N, i/N]$, and we calculate $a_i = \inf_{z \in C_i} (z + \sigma^2 \pi'(z)/2\pi(z))$ and $b_i = \sup_{z \in C_i} (z + \sigma^2 \pi'(z)/2\pi(z))$. In other words, the interval $[a_i, b_i]$ must contain the means of the distributions $p(z, \cdot)$ for all $z \in C_i$. For the proposal $\Theta_z(x, \cdot)$, let us take (independently of x) the density proportional to

$$\theta_i(y) = \begin{cases} \exp[-|y - a_i|^2 / 2\sigma^2], & \text{if } y \leq a_i, \\ 1, & \text{if } y \in [a_i, b_i], \\ \exp[-|y - b_i|^2 / 2\sigma^2], & \text{if } y \geq b_i. \end{cases}$$

It is immediately seen that $\tilde{p}(z, \cdot) \leq \theta_i(\cdot)$ for all $z \in C_i$, and this inequality is sharp. Also, take $\epsilon_i = 1$ and

$$v_i(y) = \begin{cases} \exp(-|y - b_i|^2 / 2\sigma^2), & \text{if } y \leq (a_i + b_i)/2, \\ \exp(-|y - a_i|^2 / 2\sigma^2), & \text{if } y \geq (a_i + b_i)/2, \end{cases}$$

which completes the series of estimates (12). We choose Y uniformly in $\{1, \dots, N\}$, or according to some other distribution adapted to the gradient $\pi'(z)/\pi(z)$ if desired. Since

$$v_i(y)/\theta_i(y) = \begin{cases} \exp(1/2\sigma^2)(a_i^2 - b_i^2 + 2y(b_i - a_i)), & \text{if } y \leq a_i, \\ v_i(y), & \text{if } y \in [a_i, b_i], \\ \exp(1/2\sigma^2)(b_i^2 - a_i^2 + 2y(a_i - b_i)), & \text{if } y \geq b_i, \end{cases}$$

the probability $v_i := \mathbb{P}(Y \text{ falls in } C_i \text{ and } v_i(\Phi)/\theta_i(\Phi) > \xi)$ is easy to calculate, and then we can explicitly write down the distribution of T' , the time until termination of the field coupling procedure FC-1, ..., FC-4 (Algorithm 1). This could then be optimized in N , the total size of the covering C_1, \dots, C_N .

A similar analysis can be performed for any random field on \mathbb{R}^d of the form

$$\mathfrak{X}(z) = \left(\mathbf{b}_i(z) + \sum_{j=1}^d \sigma_{ij}(z)W_j, i = 1, \dots, d \right),$$

where $\mathbf{b}(z)$ is a vector field on \mathbb{R}^d , σ is a matrix-valued field on \mathbb{R}^d and (W_1, \dots, W_d) is a standard d -dimensional Gaussian random variable. Such fields arise in the recent work by Stramer and Tweedie [12], as examples of natural uniformly ergodic Metropolis–Hastings proposals on \mathbb{R}^d .

Metropolis–Hastings-type chains comprise a class of Markov chains whose transition kernels do not possess a density (unless the state space is countable). Their general form is (see [11])

$$P(x, dy) = q(x, y) \left\{ 1 \wedge \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} \right\} \mu(dy) + r(x)\delta_x(dy), \tag{21}$$

where $r(x)$ is chosen so that $P(x, E) = 1$ and $q(x, y)\mu(dy)$ is the transition density of some arbitrary proposal chain. While it is possible to modify the field couplers to work on transition kernels of this form (this requires the proposal $\Theta_z(x, \cdot) \gg P(x, \cdot)$), it is usually very impractical to evaluate the rejection probability $r(x)$, and that renders this method inapplicable.

Fortunately, however, it is possible to proceed otherwise, as the proposal kernel $q(x, y)\mu(dy)$ is often analytically tractable. Common examples of such kernels were given as examples above, and the Langevin proposal field was explicitly analysed on a simple region.

Suppose now that we have constructed a random field $\Omega(z) \sim q(z, \cdot)\mu(\cdot)$. We shall have $\mathfrak{X}(z) \sim P(z, \cdot)$ with P given by (21) as soon as we set

$$\mathfrak{X}(z) = \begin{cases} \Omega(z), & \text{if } \pi(\Omega(z))q(\Omega(z), z) > \xi \cdot \pi(z)q(z, \Omega(z)), \\ z, & \text{otherwise,} \end{cases} \tag{22}$$

where $\xi \sim U[0, 1]$. We may view the Metropolis–Hastings accept/reject step as an operator on fields, yielding \mathfrak{X} when applied to Ω . As such, it has some useful properties.

Suppose that we generate an IID sequence of random fields $\Omega^{(1)}, \Omega^{(2)}, \dots$, and then construct $\mathfrak{X}^{(1)}, \mathfrak{X}^{(2)}, \dots$ by (22), using these to generate a flow $F_{s,t}(z) = \mathfrak{X}^{(t-1)} \circ \dots \circ \mathfrak{X}^{(s)}(z)$ for the Markov chain with transition kernel P given by (21). The random ('rejection') sets $\{z : \mathfrak{X}^{(k)}(z) = z\} \subset Z$ are not usually empty, at least with high probability.

PROPOSITION 9. *Suppose that $\Omega^{(1)}, \Omega^{(2)}, \dots$ are finitely coupled proposal fields, and define $\mathfrak{X}^{(1)}, \mathfrak{X}^{(2)}, \dots$ as a Metropolis–Hastings proposal, by (22). Then the random sets*

$$\{z : F_{0,t}(z) = z\} = \{F_t \text{ has finite range}\}^c$$

are strictly decreasing as functions of t . If the state space E is irreducible and coverable by a finite number of 1-small sets, then in fact $\{z : F_t(z) = z\} \downarrow \emptyset$, and the flow collapses in a finite time.

Proof. Since

$$\{z : F_{0,t}(z) = z\} = \{z : \mathfrak{X}^{(0)}(z) = z\} \cap \dots \cap \{z : \mathfrak{X}^{(t-1)}(z) = z\},$$

the sets are clearly decreasing. Moreover, for each $z \in E$, the Markov chain $\mathfrak{X}^{(k)}(z)$ must move eventually, and this occurs simultaneously for all z within a small set in at most geometrically distributed time. Since E is coverable by a finite number of such sets, we get $\{z : F_t(z) = z\} = \emptyset$ in finite time. After that, the map $F_t(\cdot)$ has only a finite range, and thus by Theorem 6 we get collapse eventually. \square

The result in Proposition 9 is useful for coupling procedures due to the identity $\{z : F_t(z) = z\}^c = \{F_t \text{ has finite range}\}$. This set must be monitored if we wish to detect whether $F_{0,t}(\cdot)$ has collapsed.

Note that if we do *not* assume that $\Omega^{(k)}$ is finitely coupled (only to within ϵ say), the monotonicity of the sets $\{z : F_{0,t}(z) = z\}$ is preserved, but the sets $\{F_t \text{ has finite range}\}$ are no longer their complements. Thus they can grow and shrink in complicated ways, which can make them untrackable for practical purposes.

6. Towards perfect simulation with Metropolis–Hastings chains

We have seen in the last section that if P defined by (21) is uniformly ergodic, then we can *numerically* construct a finite-time coupling of all Markov chains $X_t(x)$ with transition kernel P and deterministic initial condition $X_0(x) = x \in E$. The exact dynamics of the coupling are no longer tractable, since in fact they are random. Instead, we gain an automatic method of coupling chains, which does not depend on our ability to find clever ways to couple the flows *analytically*.

For uniformly ergodic Metropolis–Hastings chains, we describe in Algorithm 2 how to simulate a given target distribution perfectly, using the field couplers described in this paper. A large class of uniformly ergodic Metropolis–Hastings chains compatible with this method can be found in [12]. We do not discuss ways of implementing Algorithm 2 efficiently, although it is clearly possible to exploit various recurrence relations.

Nothing more needs to be said of steps (PMH-1) and (PMH-2). For step (PMH-3), it is necessary to perform some analysis, determining a lower bound

$$\inf_{z \in C_i} \frac{\pi(y)q(y, z)}{\pi(z)q(z, y)} \geq b_i(y), \quad i = 1, \dots, N, \tag{23}$$

so that we can perform the test (24) conveniently. This can typically be done by finding bounds $\pi(y)/q(z, y) \geq b'_i(y)$ and $\pi(z)/q(y, z) \leq b''_i(y)$, taking $b_i(y) = b'_i(y)/b''_i(y)$.

The step (PMH-4) can be performed efficiently if we construct tables of the correspondences $C_j \mapsto \mathfrak{X}_1 \circ \dots \circ \mathfrak{X}_n(C_j)$.

(PMH-1):

Choose a covering of E by small sets C_1, \dots, C_N and calculate the estimates (12).

(PMH-2):

Using the steps (FC-1, ..., FC-4), generate finitely coupled fields $(\Omega^{(k)} : k = -1, -2, \dots)$. For every Ω_k , keep track of the (finite) range of possible values over each C_j , $j = 1, \dots, N$.

(PMH-3):

Construct $\mathfrak{X}^{(k)}$ from $\Omega^{(k)}$ using formula (22) (with $\xi = \xi^{(k)}$), and for each C_j , given the range of values $\Omega^{(k)}(C_j) = \{q_1, \dots, q_m\}$ say, test whether

$$\pi(q_r)q(q_r, z) > \xi^{(k)}\pi(z)q(z, q_r) \quad \text{for all } z \in C_j, r = 1, \dots, m. \quad (24)$$

Let B_k denote the union of all sets C_j for which the test (24) fails. Clearly, $\{z : \mathfrak{X}^{(k)}(z) = z\} \subset B_k$.

(PMH-4):

Let $k < 0$ be given. Check that

$$B_k \cap B_{k+1} \cap \dots \cap B_1 = \emptyset.$$

If this holds, then there will typically be $k < \ell < 0$ such that $B_k \cap B_{k+1} \cap \dots \cap B_\ell = \emptyset$. For each $k \leq n \leq \ell$, let $\{q_1, \dots, q_r\} = \Omega_n(E)$ denote the (finite) range of $\Omega^{(n)}$. Compute the set

$$H_n = \{\mathfrak{X}^{(1)} \circ \dots \circ \mathfrak{X}^{(n+1)}(q_i) : i = 1, \dots, q_r\}.$$

If $H_k = H_{k+1} = \dots = H_\ell = \{p\}$, say, then the value of p is a perfect draw from π . Stop. Otherwise, try (PMH-4) with an earlier value of k .

Algorithm 2: Perfect Metropolis–Hastings

EXAMPLE 5 (LANGEVIN ALGORITHM). We consider again the Langevin random field (19) on $Z = [0, 1]$, this time considering it as a proposal field for the Metropolis–Hastings algorithm with target distribution π . Hence, up to normalization, we have $q(y, z) = \tilde{p}(y, z)$, the latter being defined by (20). To implement Algorithm 2, we need only to compute the bounds (23). Recall that we took $C_i = [(i - 1)/N, i/N]$. Let $f(z) = z + \sigma^2 \pi'(z)/2\pi(z)$ denote the mean of the density $q(z, \cdot)$. We have $f(z) \in [a_i, b_i]$ for all $z \in C_i$; consequently,

$$\begin{aligned} \pi(y)/q(z, y) &= \pi(y) \exp(|z - f(y)|^2 / 2\sigma^2), \\ &\geq \begin{cases} \pi(y), & \text{if } (i - 1)/N \leq f(y) \leq i/N, \\ \pi(y) \min_{j=i, i-1} \exp(|j/N - f(y)|^2 / 2\sigma^2), & \text{otherwise;} \end{cases} \\ \pi(z)/q(y, z) &= \pi(z) \exp(|y - f(z)|^2 / 2\sigma^2) \\ &\leq \sup_{z \in C_i} \pi(z) \exp((|y - a_i|^2 \vee |y - b_i|^2) / 2\sigma^2). \end{aligned}$$

Note that it does not matter whether we know the normalizing constant of π or not, since we are now taking the ratio of these two quantities to get $b_i(y)$. In particular, these calculations can be done on any conveniently chosen multiple of π . Setting

$$\hat{\pi}_i(y) = \frac{\pi(y)}{\sup_{z \in C_i} \pi(z)},$$

we find that

$$b_i(y) = \begin{cases} \hat{\pi}_i(y) \exp(-|y - a_i|^2 \vee |y - b_i|^2), & \text{if } (i - 1)/N \leq f(y) \leq i/N, \\ \hat{\pi}_i(y) \exp(\min_{j=i, i-1} |j/N - f(y)|^2 / 2\sigma^2 \\ \quad - |y - a_i| \vee |y - b_i|^2), & \text{otherwise.} \end{cases}$$

Acknowledgements. The authors would like to gratefully acknowledge a very thorough referee for constructive comments, and Jeff Rosenthal for informal discussions. This research was supported by the European Union TMR network ERB-FMRX-CT96-0095.

Appendix A. On-line demonstration

A Java applet that illustrates the field coupler in action in various simple examples can be viewed at:

<http://www.lbreyer.com/fcoupler.html>.

References

1. A. A. BOROVKOV and S. G. FOSS, ‘Stochastically recursive sequences and their generalizations’, *Siberian Adv. Math.* 2 (1992) 16–81. 78, 86
2. P. DIACONIS and J. A. FILL, ‘Strong stationary times via a new form of duality’, *Ann. Probab.* 18 (1990) 1483–1522. 84
3. J. A. FILL, ‘An interruptible algorithm for perfect sampling via Markov chains’, *Ann. Appl. Probab.* 8 (1998) 131–162. 86
4. S. FOSS and R. L. TWEEDIE, ‘Perfect simulation and backward coupling’, *Comm. Statist. Stochastic Models* (special issue in honor of Marcel F. Neuts) 14 (1998) 187–203. 88
5. P. J. GREEN and D. J. MURDOCH ‘Exact sampling for Bayesian inference: towards general-purpose algorithms’, *Bayesian statistics IV* (ed. J. Bernardo, J. Berger, A. P. Dawid and A. F. M. Smith, Oxford University Press, 1999) 301–321.
6. T. LINDVALL, *Lectures in the coupling method* (John Wiley & Sons, New York, 1992). 77
7. S. P. MEYN and R. L. TWEEDIE, *Markov chains and stochastic stability* (Springer, London, 1993). 77
8. D. J. MURDOCH and P. J. GREEN, ‘Exact sampling from a continuous state space’, *Scand. J. Statist.* 25 (1998) 483–502. 89
9. J. PROPP and D. B. WILSON, ‘Exact sampling with coupled Markov chains and applications to statistical mechanics’, *Random Structures Algorithms* 9 (1996) 223–252. 78, 78, 84
10. G. O. ROBERTS and J. S. ROSENTHAL, ‘Quantitative bounds for convergence rates of continuous time Markov processes’, *Electron. J. Probab.* 1 (1996). 80

11. G. O. ROBERTS and A. F. M. SMITH, 'Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods', *J. Roy. Statist. Soc. Ser. B* 55 (1993) 3–23. 79, 90
12. O. STRAMER and R. L. TWEEDIE, 'Self-targeting candidates for Metropolis–Hastings algorithms', *Methodology and Computing in Applied Probability* 1 (1999) 307–328. 82, 90, 91

L. A. Breyer laird@lbreyer.com

G. O. Roberts g.o.roberts@lancaster.ac.uk

Department of Mathematics and Statistics

Lancaster University

Lancaster LA1 4YF