

Partially Collapsed Gibbs Samplers: Theory and Methods

David A. VAN DYK and Taeyoung PARK

Ever-increasing computational power, along with ever-more sophisticated statistical computing techniques, is making it possible to fit ever-more complex statistical models. Among the more computationally intensive methods, the Gibbs sampler is popular because of its simplicity and power to effectively generate samples from a high-dimensional probability distribution. Despite its simple implementation and description, however, the Gibbs sampler is criticized for its sometimes slow convergence, especially when it is used to fit highly structured complex models. Here we present partially collapsed Gibbs sampling strategies that improve the convergence by capitalizing on a set of functionally incompatible conditional distributions. Such incompatibility generally is avoided in the construction of a Gibbs sampler, because the resulting convergence properties are not well understood. We introduce three basic tools (marginalization, permutation, and trimming) that allow us to transform a Gibbs sampler into a partially collapsed Gibbs sampler with known stationary distribution and faster convergence.

KEY WORDS: Blocking; Gibbs sampler; Incompatible Gibbs sampler; Marginal data augmentation; Rate of convergence.

1. INTRODUCTION

The *ordinary Gibbs sampler* (Geman and Geman 1984) begins with the joint posterior distribution of a set of unknown quantities and updates groups of these quantities by sampling them from their conditional distributions under the joint posterior distribution. Each quantity is generally updated exactly once in each iteration. The partially collapsed Gibbs (PCG) sampler replaces some of these conditional distributions with conditional distributions under some *marginal distributions* of the joint posterior distribution. This strategy is useful because it can result in samplers with much better convergence properties, and it is interesting because it may result in sampling from a set of *incompatible* conditional distributions; that is, there may be no joint distribution that corresponds to this set of conditional distributions.

Our technique can be viewed as a generalization of blocking (Liu, Wong, and Kong 1994) in that the resulting conditional distributions sometimes can be combined to form a Gibbs sampler that is a blocked version of the original sampler. In such cases, we can recover a set of compatible conditional distributions and an ordinary Gibbs sampler by combining steps. This is not always possible, however, and some PCG samplers may be composed only of draws from incompatible conditional distributions. In this regard, PCG samplers constitute a generalization of the Gibbs sampler, in that ordinary Gibbs samplers are constructed using the conditional distributions of some joint distribution. Because both blocking and collapsing (Liu et al. 1994) are special cases of PCG, our methods can be viewed as a generalizing and unifying framework for these important and efficient methods. Like blocked and collapsed samplers, PCG samplers dominate their parent Gibbs samplers in terms of their convergence and maintain the target posterior distribution as their stationary distribution.

To transform a Gibbs sampler into a PCG sampler, we use three basic tools. The first tool, *marginalization*, entails moving a group of unknowns from being conditioned upon to being

sampled in one or more steps of a Gibbs sampler; the marginalized group can differ among the steps. Second, we may need to *permute* the steps of the sampler to use the third tool, *trimming* sampled components from the various steps that can be removed from the sampler without altering its Markov transition kernel. Marginalization and permutation both trivially maintain the stationary distribution of a Gibbs sampler, and both can effect the convergence properties of the chain; marginalization can dramatically improve convergence, whereas the effect of a permutation is typically small. Trimming, on the other hand, is explicitly designed to maintain the kernel of the chain. Its primary advantage is to reduce the complexity and the computational burden of the individual steps. It is trimming that introduces incompatibility into the sampler.

Although our methods are motivated by computational challenges in applied problems in astronomy and multiple imputation, to save space here we focus on our methods and their properties. Interested readers can find detailed applications in our technical paper (Park and van Dyk 2008).

2. MOTIVATING EXAMPLES

To illustrate PCG samplers, we consider the simple random-effects model given by

$$y_{ij} = \xi_i + \varepsilon_{ij}, \quad \text{for } i = 1, \dots, k \text{ and } j = 1, \dots, n, \quad (1)$$

where $\xi_i \stackrel{\text{iid}}{\sim} N(\mu, \tau^2)$ and $\varepsilon_{ij} \stackrel{\text{iid}}{\sim} N(0, \sigma^2)$ are independent with y_{ij} observation j in group i , n the number of units in each group, ξ_i the mean of group i , μ the mean of the group means, τ^2 the between group variance, σ^2 the within-group variance, and τ^2 and σ^2 presumed known. From a Bayesian perspective, we are interested in the joint posterior distribution $p(\xi, \mu | \mathbf{Y})$ computed under the flat prior distribution $p(\mu) \propto 1$, where $\xi = (\xi_1, \xi_2, \dots, \xi_k)$ and $\mathbf{Y} = \{y_{ij}, i = 1, \dots, k, j = 1, \dots, n\}$. To fit this random-effects model, we can use a prototype two-step Gibbs sampler that iterates between the following two steps:

Sampler 1

Step 1. Draw $\xi^{(t)}$ from $p(\xi | \mu^{(t-1)}, \mathbf{Y})$,

David A. van Dyk is Professor, Department of Statistics, University of California Irvine, Irvine, CA 92697 (E-mail: dvd@ics.uci.edu). Taeyoung Park is Assistant Professor, Department of Statistics, University of Pittsburgh, Pittsburgh, PA 15260 (E-mail: tpark@pitt.edu). This project was supported in part by National Science Foundation grant DMS-01-04129, DMS-04-38240, and DMS-04-06085 and by National Aeronautic and Space Administration contracts NAS8-39073 and NAS8-03060 (CXC).

where

$$\xi_i | \mu^{(t-1)}, \mathbf{Y} \stackrel{\text{ind}}{\sim} \mathbf{N} \left(\frac{n\tau^2 \bar{Y}_i + \sigma^2 \mu^{(t-1)}}{n\tau^2 + \sigma^2}, \frac{\tau^2 \sigma^2}{n\tau^2 + \sigma^2} \right)$$

for $i = 1, \dots, k$.

Step 2. Draw $\mu^{(t)}$ from

$$p(\mu | \boldsymbol{\xi}^{(t)}, \mathbf{Y}) = \mathbf{N} \left(\frac{\sum_i \xi_i^{(t)}}{k}, \frac{\tau^2}{k} \right),$$

with $\bar{Y}_i = \sum_{j=1}^n Y_{ij}/n$ being the mean of the observations in group i . We emphasize that this is a toy example introduced for illustrative purposes; there is no need for Gibbs sampling when both the between and within group variances are known.

In the case of a two-step Gibbs sampler, the geometric rate of convergence is the same as the lag-one autocorrelation; in Sampler 1, the lag-one autocorrelation is the shrinkage parameter, $\sigma^2/(n\tau^2 + \sigma^2)$. To illustrate the potentially slow convergence of Sampler 1, we simulate test data with $\mu = 0$, $\sigma = 10$, and $\tau = .1$ when there are $n = 10$ observations in each of $k = 10$ groups. The test data are fit by running Sampler 1 with 5,000 iterations; the resulting lag-one autocorrelation for μ is .999. The first row of Figure 1 presents the output of Sampler 1 and illustrates the poor mixing and high autocorrelations of the subchain for μ and the strong posterior correlation of μ and ξ_1 .

To improve the convergence of a Markov chain constructed with a Gibbs sampler, we can replace a conditional distribution of the original Gibbs sampler with a conditional distribution of a *marginal distribution* of the target distribution. Throughout this article, such a conditional distribution that conditions on fewer unknown components is referred to as a *reduced conditional distribution*. Reduced conditional distributions are conditional distributions of a marginal distribution of the target joint distribution. In the random-effects model, we consider the marginal distribution $p(\mu | \mathbf{Y}) = \int p(\boldsymbol{\xi}, \mu | \mathbf{Y}) d\boldsymbol{\xi}$ of the target distribution $p(\boldsymbol{\xi}, \mu | \mathbf{Y})$. We replace step 2 of Sampler 1 with the trivial “conditional” distribution of this marginal distribution. This substitution yields the following sampler:

Sampler 2

Step 1. Draw $\boldsymbol{\xi}^{(t)}$ from $p(\boldsymbol{\xi} | \mu^{(t-1)}, \mathbf{Y})$.

Step 2. Draw $\mu^{(t)}$ from

$$p(\mu | \mathbf{Y}) = \mathbf{N} \left(\frac{\sum_i \sum_j y_{ij}}{nk}, \frac{n\tau^2 + \sigma^2}{nk} \right).$$

Here step 2 simulates μ directly from its marginal posterior distribution. The advantage of this strategy is clear: We immediately obtain independent draws of μ from the target posterior distribution. But the two conditional distributions used in Sampler 2, $p(\boldsymbol{\xi} | \mu, \mathbf{Y})$ and $p(\mu | \mathbf{Y})$, are incompatible and imply inconsistent dependence structure. Even in this simple case, the incompatible conditional distributions improve the convergence of the sampler, but at the expense of the correlation structure of the target distribution. Indeed, because $\mu^{(t)}$ is sampled independently of $\boldsymbol{\xi}^{(t)}$, the Markov chain has stationary distribution $p(\boldsymbol{\xi} | \mathbf{Y})p(\mu | \mathbf{Y})$ rather than $p(\boldsymbol{\xi}, \mu | \mathbf{Y})$. This is illustrated in the second row of Figure 1, which confirms that the subchain for μ converges immediately to its target marginal distribution, but the correlation structure between μ and ξ_1 (and all of $\boldsymbol{\xi}$) is lost.

There is an obvious solution. Sampler 2 first draws $\boldsymbol{\xi}$ from its conditional posterior distribution $p(\boldsymbol{\xi} | \mu, \mathbf{Y})$ and then draws μ from its marginal posterior distribution $p(\mu | \mathbf{Y})$, rather than vice versa. By simply exchanging the order of the steps, we regain the correlation structure of the target distribution. The resulting Gibbs sampler iterates between two steps:

Sampler 3

Step 1. Draw $\mu^{(t)}$ from $p(\mu | \mathbf{Y})$.

Step 2. Draw $\boldsymbol{\xi}^{(t)}$ from $p(\boldsymbol{\xi} | \mu^{(t)}, \mathbf{Y})$.

Sampler 3 comprises two incompatible conditional distributions and converges quicker than Sampler 1, while maintaining the correlations of the target distribution. Of course in this case, the PCG sampler (Sampler 3) is simply a blocked version of Sampler 1; steps 1 and 2 combine into a single independent draw from the target distribution. As we show, however, PCG samplers can be more general than blocked Gibbs samplers when there are more than two steps. The bottom row of Figure 1 illustrates the fast convergence of the subchain for μ and the correct correlation structure of μ and ξ_1 .

We now consider a more complex four-step prototype Gibbs sampler with target distribution $p(\mathbf{W}, \mathbf{X}, \mathbf{Y}, \mathbf{Z})$. As the number of components in a Gibbs sampler increases, so do the ways to construct PCG samplers; here we focus on an example in which partial collapse does not correspond with blocking. Generally this situation is even more complicated when the sampled component are vectors, in that we may marginalize out certain subvectors; see Park and van Dyk (2008). We begin with the Gibbs sampler that iterates among the following steps:

Sampler 4

Step 1. Draw \mathbf{W} from $p(\mathbf{W} | \mathbf{X}, \mathbf{Y}, \mathbf{Z})$.

Step 2. Draw \mathbf{X} from $p(\mathbf{X} | \mathbf{W}, \mathbf{Y}, \mathbf{Z})$.

Step 3. Draw \mathbf{Y} from $p(\mathbf{Y} | \mathbf{W}, \mathbf{X}, \mathbf{Z})$.

Step 4. Draw \mathbf{Z} from $p(\mathbf{Z} | \mathbf{W}, \mathbf{X}, \mathbf{Y})$.

Suppose that it is possible to directly sample from $p(\mathbf{Y} | \mathbf{X}, \mathbf{Z})$ and $p(\mathbf{Z} | \mathbf{X}, \mathbf{Y})$, which are both conditional distributions of $\int p(\mathbf{W}, \mathbf{X}, \mathbf{Y}, \mathbf{Z}) d\mathbf{W}$. By replacing steps 3 and 4 with draws from these two distributions, we are partially collapsing \mathbf{W} out of Sampler 4. But substituting the conditional distributions of a marginal distribution of the target distribution into a Gibbs sampler may result in a transition kernel with unknown stationary distribution. This is illustrated by the loss of correlation structure that occurs when using Sampler 2; see the last column of Figure 1. Nevertheless, we hope to capitalize on the potential computational gain that partial collapse offers. Thus our goal is to formalize a procedure that allows us to introduce partially collapsed steps while ensuring that the target stationary distribution is maintained. We illustrate our strategy in this example and formalize it in Section 3.

Moving components in a step of a Gibbs sampler from being conditioned on to being sampled can improve the convergence characteristics of the sampler. This neither alters the stationary distribution of the chain nor destroys the compatibility of the conditional distributions. For example, based on the available reduced conditional distributions, we can sample \mathbf{W} jointly with \mathbf{Y} in step 3 and with \mathbf{Z} in step 4; for example, \mathbf{W} and \mathbf{Y} can be sampled jointly by first sampling from $P(\mathbf{Y} | \mathbf{X}, \mathbf{Z})$ and then from $P(\mathbf{W} | \mathbf{Y}, \mathbf{X}, \mathbf{Z})$, both of which are assumed to be tractable. The resulting Gibbs sampler iterates among the following steps:

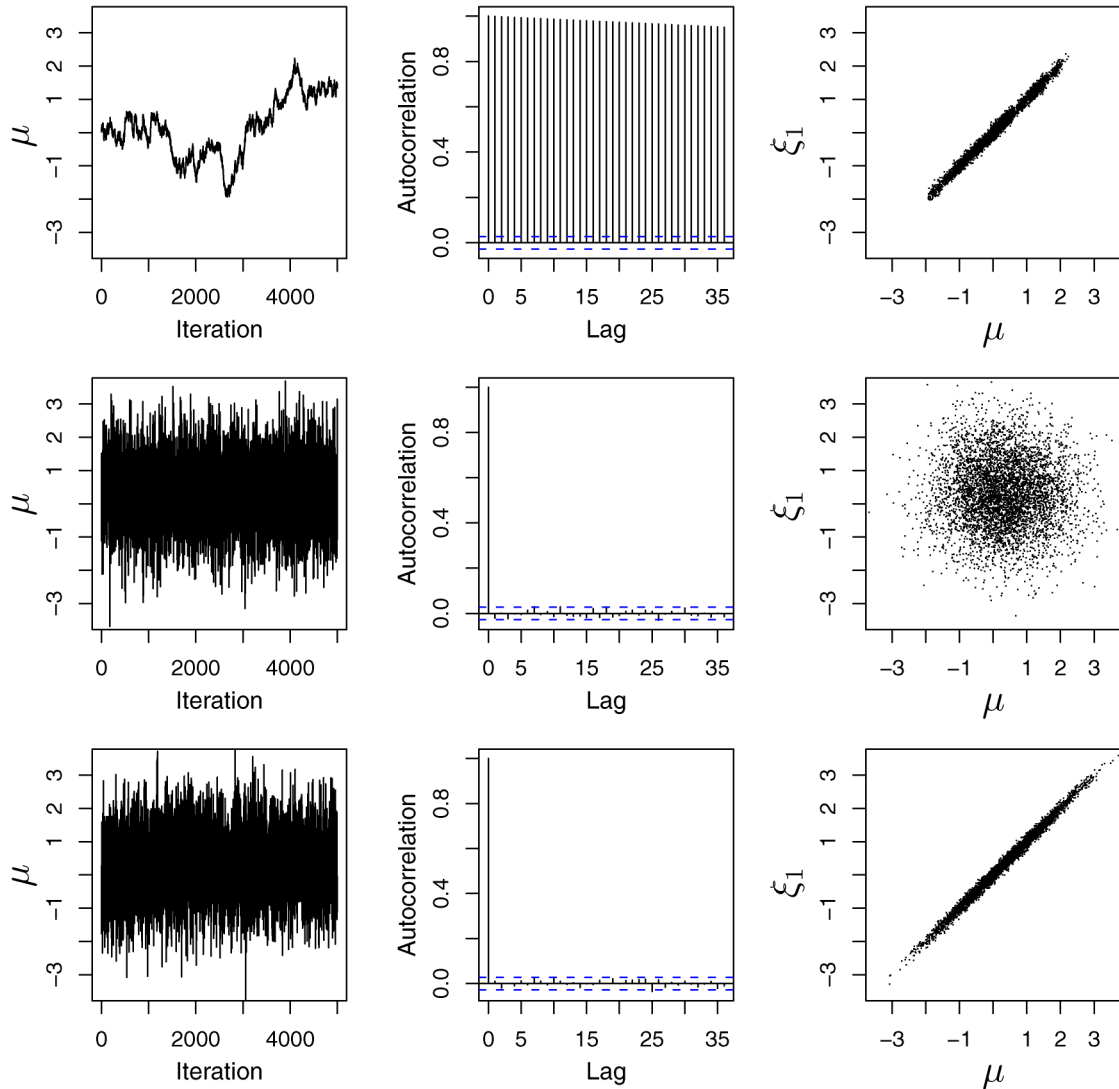


Figure 1. Comparison of three samplers for the simple random-effects model. The first two columns show the mixing and autocorrelations of the subchain for μ ; the last column shows the correlation structure between μ and ξ_1 . The three rows represent the ordinary Gibbs sampler (sampler 1), the Gibbs sampler resulting from the inappropriate substitution of a reduced conditional distribution (sampler 2), and the PCG sampler (sampler 3).

Sampler 5

- Step 1. Draw \mathbf{W}^* from $p(\mathbf{W}|\mathbf{X}, \mathbf{Y}, \mathbf{Z})$.
- Step 2. Draw \mathbf{X} from $p(\mathbf{X}|\mathbf{W}, \mathbf{Y}, \mathbf{Z})$.
- Step 3. Draw $(\mathbf{W}^*, \mathbf{Y})$ from $p(\mathbf{W}, \mathbf{Y}|\mathbf{X}, \mathbf{Z})$.
- Step 4. Draw (\mathbf{W}, \mathbf{Z}) from $p(\mathbf{W}, \mathbf{Z}|\mathbf{X}, \mathbf{Y})$.

In each step we condition on the most recently sampled value of each variable that is not sampled in that step. Thus in step 2, we condition on the $\mathbf{W} = \mathbf{W}^*$ drawn in step 1. The output comprises the most recently sampled values of each variable at the end of the iteration, that is, \mathbf{W} drawn in step 4, \mathbf{X} drawn in step 2, \mathbf{Y} drawn in step 3, and \mathbf{Z} drawn in step 4. Here and elsewhere a superscript “*” designates an *intermediate quantity* that is sampled but is not part of the output of an iteration. Sampler 5 is a trivial generalization of what is typically considered an ordinary Gibbs sampler, with \mathbf{W} sampled more than once during an iteration. This sampler may be inefficient, in that it draws \mathbf{W} three times in each iteration. But removing any two draws from the iteration necessarily affects the transition kernel, because the first draw is conditioned on in the next step and

the third draw is part of the output of the sampler. As Figure 1 illustrates, such changes to the transition kernel can destroy the correlation structure of the stationary distribution or otherwise affect convergence of the chain.

In general, we consider removing only draws of intermediate quantities from a sampler, because removing draws of any part of the output quantities and replacing output quantities with corresponding intermediate quantities necessarily alters the transition kernel and may affect the stationary distribution. Moreover, we only remove draws of intermediate quantities if removing them from the iteration does not affect the transition kernel. Permuting the steps of a Gibbs sampler does not alter its stationary distribution, but sometimes allows us to meet the criteria for removing redundant draws. In the case of Sampler 5, such permutation yields a Gibbs sampler that iterates among the following:

Sampler 6

- Step 1. Draw $(\mathbf{W}^*, \mathbf{Y})$ from $p(\mathbf{W}, \mathbf{Y}|\mathbf{X}, \mathbf{Z})$.

- Step 2. Draw $(\mathbf{W}^*, \mathbf{Z})$ from $p(\mathbf{W}, \mathbf{Z}|\mathbf{X}, \mathbf{Y})$.
 Step 3. Draw \mathbf{W} from $p(\mathbf{W}|\mathbf{X}, \mathbf{Y}, \mathbf{Z})$.
 Step 4. Draw \mathbf{X} from $p(\mathbf{X}|\mathbf{W}, \mathbf{Y}, \mathbf{Z})$.

Here the first two draws of \mathbf{W} correspond to intermediate quantities that are not conditioned on and are not part of the output. This permutation alters the transition kernel while maintaining the stationary distribution, and allows us to remove the two redundant draws of \mathbf{W} without changing the transition kernel. Removing the intermediate quantities \mathbf{W}^* from Sampler 6 yields a PCG sampler that iterates among the following:

Sampler 7

- Step 1. Draw \mathbf{Y} from $p(\mathbf{Y}|\mathbf{X}, \mathbf{Z})$.
 Step 2. Draw \mathbf{Z} from $p(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$.
 Step 3. Draw \mathbf{W} from $p(\mathbf{W}|\mathbf{X}, \mathbf{Y}, \mathbf{Z})$.
 Step 4. Draw \mathbf{X} from $p(\mathbf{X}|\mathbf{W}, \mathbf{Y}, \mathbf{Z})$.

We can block steps 2 and 3 in Sampler 7 into a joint draw from $p(\mathbf{W}, \mathbf{Z}|\mathbf{X}, \mathbf{Y})$, thereby yielding the following:

Sampler 8

- Step 1. Draw \mathbf{Y} from $p(\mathbf{Y}|\mathbf{X}, \mathbf{Z})$.
 Step 2. Draw \mathbf{W} from $p(\mathbf{W}, \mathbf{Z}|\mathbf{X}, \mathbf{Y})$.
 Step 3. Draw \mathbf{X} from $p(\mathbf{X}|\mathbf{W}, \mathbf{Y}, \mathbf{Z})$.

The three conditional distributions in Sampler 8 remain incompatible, however. Thus this PCG sampler does not simply correspond to a blocked version of Sampler 4. This illustrates that partial collapse is a more general technique than blocking.

The resulting PCG sampler (i.e., Sampler 8) is not an ordinary Gibbs sampler; permuting its draws may result in a kernel with unknown stationary distribution. As in this case, however, we can sometimes verify that a PCG sampler is valid in that the resulting stationary distribution is equal to the target distribution. Still, because the removal of intermediate quantities introduces incompatibility, removal must be done with great care.

3. BASIC TOOLS

Here we present three basic tools for constructing PCG samplers. Unless marginalized quantities are removed from the iteration with care, the resulting chain may not converge properly. Thus the tools are designed to ensure that the resulting PCG samplers converge quickly to the target distribution. We discuss the three tools—marginalization, permutation, and trimming—in the order that they are applied. Of these, trimming must be done with care, because it has the potential to alter the chain's stationary distribution. Figure 2 presents a flowchart delineating how the basic tools are applied to an ordinary Gibbs sampler to construct a PCG sampler with the same stationary distribution. In the following sections, we closely examine each component of the flowchart in Figure 2.

3.1 Marginalization

We aim to construct a PCG sampler with stationary distribution $p(\mathbf{X})$, where \mathbf{X} is a vector quantity that we partition into J subvectors, $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_J)$. Consider the sequence of index sets $\mathcal{J} = \{\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_P\}$, where $\mathcal{J}_p \subset \{1, 2, \dots, J\}$ for $p = 1, 2, \dots, P$ such that $\bigcup_{p=1}^P \mathcal{J}_p = \{1, 2, \dots, J\}$. Let \mathcal{J}_p^c be the complement of \mathcal{J}_p in $\{1, 2, \dots, J\}$ and $\mathcal{X}_{\mathcal{J}_p}$ be the collection of components of \mathbf{X} corresponding to the index set \mathcal{J}_p , that is, $\mathcal{X}_{\mathcal{J}_p} = \{\mathbf{X}_j : j \in \mathcal{J}_p\}$. Step p of a P -step Gibbs sampler can be written as follows:

- Step p . Draw $\mathcal{X}_{\mathcal{J}_p}^{(t)}$ from $p(\mathcal{X}_{\mathcal{J}_p} | \mathcal{X}_{\mathcal{J}_p^c}^{(t-1)})$, for $p = 1, 2, \dots, P$,

where $\mathcal{X}_{\mathcal{J}_p}^{(t-1)}$ consists of the most recently sampled values of each component of $\mathcal{X}_{\mathcal{J}_p}$. This is a Gibbs sampler that uses compatible conditional distributions, where some components of \mathbf{X} may be updated in multiple steps; thus it is not an ordinary Gibbs sampler.

An ordinary Gibbs sampler updates each (vector) component of \mathbf{X} only once in an iteration. In our notation this corresponds to the case where \mathcal{J} is a partition of $\{1, 2, \dots, J\}$. At the other extreme, suppose that there exists an index k such that $k \in \mathcal{J}_p$ for each p . In this case \mathbf{X}_k is drawn in each step and is never conditioned on; we say that \mathbf{X}_k has been (completely) collapsed out of the Gibbs sampler. Thus we can reformulate the Gibbs sampler in terms of the marginal distribution $\int \pi(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_J) d\mathbf{X}_k$, without altering its kernel.

The first step in constructing a PCG sampler is to marginalize some components of \mathbf{X} out of some steps of the sampler. To do this, we replace \mathcal{J}_q with $\tilde{\mathcal{J}}_q$ for one or more $q \in \{1, 2, \dots, P\}$, where \mathcal{J}_q is a proper subset of $\tilde{\mathcal{J}}_q$. That is, in step q , we move some components of \mathbf{X} from being conditioned on to being sampled. As we show later, the marginalization can improve the convergence properties of the Gibbs sampler; see Section 4 for the theory on the improved rate of convergence. Then step q is conditional on fewer components of \mathbf{X} :

- Step q . Draw $\mathcal{X}_{\tilde{\mathcal{J}}_q}^{(t)}$ from $p(\mathcal{X}_{\tilde{\mathcal{J}}_q} | \mathcal{X}_{\tilde{\mathcal{J}}_q^c}^{(t-1)})$.

Marginalizing out some components of \mathbf{X} alters the transition kernel, but not the stationary distribution for \mathbf{X} or the compatibility of the conditional distributions. The improved rate of convergence of the sampler is attributed mainly to marginalization.

3.2 Permutation

In the case of a P -step Gibbs sampler, the steps can be reordered into $P!$ possible permutations. Permuting the compatible conditional distributions of a Gibbs sampler typically changes its transition kernel and interchanges intermediate quantities with output quantities, while maintaining the stationary distribution of the chain. Our goal in permuting the steps is to arrange them so that as many of the marginalized components as possible are intermediate quantities that are not conditioned on in subsequent steps.

Permutation of the steps may affect the convergence of a sampler, but its influence is typically small compared with that of marginalization (van Dyk and Meng 1997). In this article we tend to ignore the effect of permutation on convergence; we are interested in permutations merely, because they can allow removal of some intermediate quantities.

3.3 Trimming

By trimming, we mean discarding a subset of the components that were to be sampled in one or more steps of a Gibbs sampler. In the P -step Gibbs sampler, for example, trimming the marginalized components of \mathbf{X} in step q yields the following:

- Step q . Draw $\mathcal{X}_{\mathcal{J}_q}^{(t)}$ from $p(\mathcal{X}_{\mathcal{J}_q} | \mathcal{X}_{\mathcal{J}_q^c}^{(t-1)})$.

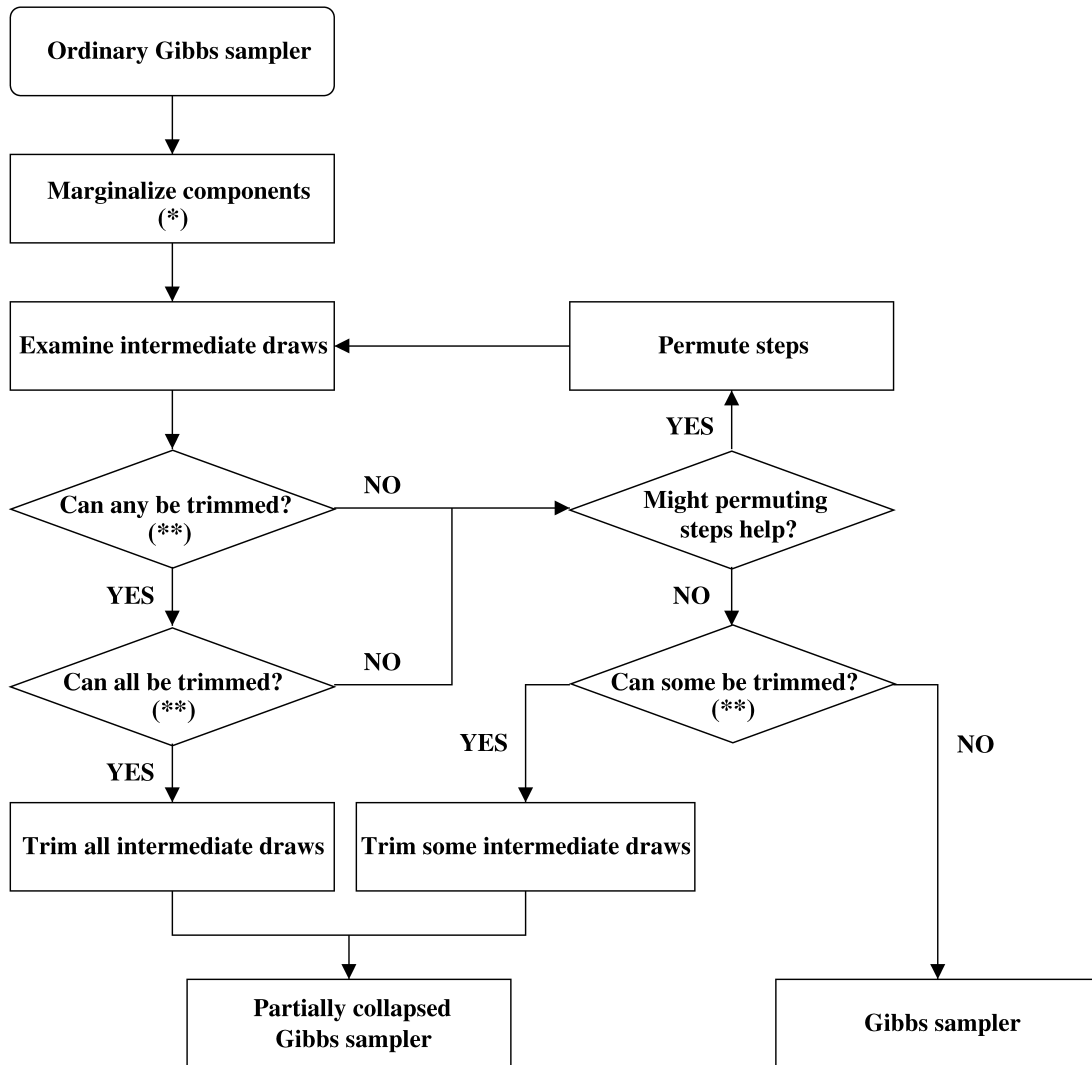


Figure 2. Flow diagram for deriving a PCG sampler from an ordinary Gibbs sampler. (*) As many components should be marginalized in as many steps as possible. (**) Intermediate draws should only be trimmed if they are not conditional on in subsequent draws.

The reduced conditional distribution sampled in this step typically is not compatible with the other conditional distributions sampled in the sampler. In particular, because $\mathcal{J}_q \cup \tilde{\mathcal{J}}_q^c$ is not equal to \mathcal{J} (i.e., \mathcal{J}_q is a proper subset of $\tilde{\mathcal{J}}_q$), this conditional distribution is not defined on the same space as the conditional distributions of the original sampler. Thus trimming introduces incompatibility into the conditional distributions of a PCG sampler. This means the resulting PCG sampler may no longer be a Gibbs sampler per se, because Gibbs samplers are generally expected to be constructed with compatible conditional distributions. Unlike in a Gibbs sampler, permuting the steps of a PCG sampler may result in a new Markov transition kernel with an unknown stationary distribution. Nonetheless, trimming is advantageous, because each iteration is less computationally demanding. Indeed, trimming may render an intractable sampling step tractable (see Park and van Dyk 2008).

We emphasize that trimming must be done carefully, because it has the potential to alter the stationary distribution of the chain. Intermediate quantities may be conditioned on in subsequent draws. Thus we can trim only intermediate quantities that are not conditioned on if we hope to maintain the transition ker-

nel. Trimming intermediate quantities that are conditioned on in subsequent steps can affect the transition kernel and the correlation structure of the stationary distribution. Thus care must be taken when trimming intermediate quantities to maintain the stationary distribution.

4. PARTIALLY COLLAPSED GIBBS THEORY

To discuss the effect of partial collapse on convergence, we introduce some technical concepts concerning Markov chains. We follow the notation of Liu (2001, sec. 6.7). Let $L^2(\pi)$ denote the set of all functions $h(\mathbf{X})$ such that $\int h^2(\mathbf{X})\pi(\mathbf{X})d\mathbf{X} < \infty$. This set is a Hilbert space with inner product $\langle h, g \rangle = E_\pi\{h(\mathbf{X})g(\mathbf{X})\}$, so that $\|h\|^2 = \text{var}_\pi(h)$. For a general Markov chain $\mathcal{M}_{\mathbf{X}} = \{\mathbf{X}^{(0)}, \mathbf{X}^{(1)}, \dots\}$ with transition kernel $\mathcal{K}(\mathbf{X}^{(1)} = \mathbf{X} | \mathbf{X}^{(0)} = \mathbf{X}')$, we define the forward operator \mathbf{F} on $L^2(\pi)$ for $\mathcal{M}_{\mathbf{X}}$ by

$$\mathbf{F}h(\mathbf{X}') = \int h(\mathbf{X})\mathcal{K}(\mathbf{X}|\mathbf{X}')d\mathbf{X} = E\{h(\mathbf{X}^{(1)})|\mathbf{X}^{(0)} = \mathbf{X}'\}.$$

Let $\mathbf{L}_0^2(\pi) = \{h : E_\pi\{h(\mathbf{X})\} = 0, \text{var}_\pi\{h(\mathbf{X})\} < \infty\}$. This is also a Hilbert space with the same inner product and is invariant

under \mathbf{F} . We define \mathbf{F}_0 to be the forward operator on $L_0^2(\pi)$ induced by \mathbf{F} . If we define the norm of this forward operator by $\|\mathbf{F}_0\| = \sup_h \|\mathbf{F}_0 h(\mathbf{X})\|$ with the supremum taken over $h \in L_0^2(\pi)$, then it can be shown that

$$\begin{aligned} \|\mathbf{F}_0\| &= \sup_{h \in L_0^2(\pi)} (\text{var}_\pi [\mathbb{E}\{h(\mathbf{X}^{(1)})|\mathbf{X}^{(0)}\}])^{1/2} \\ &= \sup_{h \in L_0^2(\pi)} \{E_\pi([\mathbb{E}\{h(\mathbf{X}^{(1)})|\mathbf{X}^{(0)}\}]^2)\}^{1/2} = \rho(\mathbf{X}^{(1)}, \mathbf{X}^{(0)}), \end{aligned}$$

where $\rho(\vartheta, \varphi)$ as the maximum correlation of ϑ and φ ,

$$\begin{aligned} \rho(\vartheta, \varphi) &= \sup \text{corr}\{h(\vartheta), g(\varphi)\} \\ &= \sup_{h: \text{var}\{h(\vartheta)\}=1} (\text{var}_\pi [\mathbb{E}\{h(\vartheta)|\varphi\}])^{1/2}, \end{aligned}$$

where the first sup is over all nonconstant scalar functions h and g with finite variance (see, e.g., Liu et al. 1994). Here the maximum autocorrelation $\rho(\mathbf{X}^{(1)}, \mathbf{X}^{(0)})$ is computed under the stationary distribution of $\mathcal{M}_{\mathbf{X}}$, and also is denoted by $\rho(\mathcal{M}_{\mathbf{X}})$.

The spectral radius of \mathbf{F}_0 , $r(\mathbf{F}_0)$, typically governs the convergence of $\mathcal{M}_{\mathbf{X}}$ (Liu 2001) and is related to the norm

$$\begin{aligned} &\text{by } \lim_{n \rightarrow \infty} \|\mathbf{F}_0^n\|^{1/n} = r(\mathbf{F}_0) \text{ and} \\ &\text{by the inequality } r(\mathbf{F}_0) \leq \|\mathbf{F}_0\|. \end{aligned} \quad (2)$$

Along with the relationship between the maximum autocorrelation of $\mathcal{M}_{\mathbf{X}}$ and $\|\mathbf{F}_0\|$, (2) justifies using $\|\mathbf{F}_0\|$ in the analysis of the convergence behavior of $\mathcal{M}_{\mathbf{X}}$.

Consider the P -step Gibbs sampler described in Section 3. We define a p -step-lagged Gibbs sampler for $p = 0, 1, \dots, P-1$, as the Gibbs sampler with iteration that begins with step $p+1$, cycles through the steps in the same order as in the original sampler, and ends with step p . The forward operators of the P p -step-lagged Gibbs samplers (with $p = 0, 1, \dots, P-1$) have the same spectral radius, which we call r ; however, they may have different norms and maximum autocorrelations. We denote the maximum autocorrelation $\rho(\mathcal{M}_{\mathbf{X}})$ of the p -step-lagged chain by ρ_p for $p = 0, 1, \dots, P-1$. By (2), we have

$$r \leq \min_{p \in \{0, 1, \dots, P-1\}} \rho_p, \quad (3)$$

where we introduce $\min_{p \in \{0, 1, \dots, P-1\}} \rho_p$ as the *cyclic-permutation bound* on the spectral radius. Later we show that by marginalizing a component of \mathbf{X} in step $p+1$ (i.e., the first step of the p -step-lagged Gibbs sampler), we reduce ρ_p , thereby reducing the bound given in (3) on the spectral radius. This leads to the following result.

Theorem 1. Sampling more components of \mathbf{X} in any set of steps of a Gibbs sampler can only reduce the cyclic-permutation bound on the spectral radius of that Gibbs sampler.

It remains only to show that marginalizing a component of \mathbf{X} in step $p+1$ reduces ρ_p . Because step $p+1$ is the first step of the p -step-lagged Gibbs sampler, we evaluate the effect of marginalizing a component of \mathbf{X} in step $p+1$ on ρ_p . This is because the following theorem evaluates the effect of marginalization in the first step of a Gibbs sampler. To illustrate the computational advantages of the partial collapse, we consider the generic P -step Gibbs sampler introduced in Section 3, from which we marginalize some

components of \mathbf{X} in step 1. We wish to compare two sequences of index sets and their resulting transition kernels; namely $(\mathcal{X}_{\mathcal{J}_1}, \mathcal{X}_{\mathcal{J}_2}, \dots, \mathcal{X}_{\mathcal{J}_P})$ and its kernel $\mathcal{K}(\mathbf{X}|\mathbf{X}')$ and $(\tilde{\mathcal{X}}_{\tilde{\mathcal{J}}_1}, \tilde{\mathcal{X}}_{\tilde{\mathcal{J}}_2}, \dots, \tilde{\mathcal{X}}_{\tilde{\mathcal{J}}_P})$ and its kernel $\tilde{\mathcal{K}}(\mathbf{X}|\mathbf{X}')$, where $\mathcal{J}_p = \tilde{\mathcal{J}}_p$ for $p = 2, \dots, P$, but $\mathcal{X}_{\mathcal{J}_1} = \{\mathbf{x}_1\}$ and $\tilde{\mathcal{X}}_{\tilde{\mathcal{J}}_1} = \{\mathbf{x}_1, \mathbf{x}_2\}$ with $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$. Here $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ is an alternate partition of $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_J)$ introduced to simplify notation in the theorem, that is, $\mathcal{J}_1 \subset \tilde{\mathcal{J}}_1 \subset \{1, 2, \dots, J\}$, where both subsets are proper subsets, $\mathbf{x}_1 = \{\mathbf{X}_j : j \in \mathcal{J}_1\}$, $\mathbf{x}_2 = \{\mathbf{X}_j : j \in \tilde{\mathcal{J}}_1 \setminus \mathcal{J}_1\}$, and $\mathbf{x}_3 = \{\mathbf{X}_j : j \in \tilde{\mathcal{J}}_1^c\}$. In words, the two sequence of index sets represent identical samplers, except in step 1, where more components of \mathbf{X} are drawn in the Gibbs sampler with kernel $\tilde{\mathcal{K}}(\mathbf{X}|\mathbf{X}')$. In this case we have the following result.

Theorem 2. Sampling more components of \mathbf{X} in the first step of a Gibbs sampler improves the resulting maximal autocorrelation, $\rho(\mathcal{M}_{\mathbf{X}})$.

Proof. Let h be an arbitrary function of \mathbf{X} with mean 0 and finite variance under stationarity, that is, $h \in L_0^2(\pi)$ with π the stationary distribution of \mathbf{X} ; then

$$\begin{aligned} \tilde{\mathbb{E}}\{h(\mathbf{X})|\mathbf{x}'_3\} &= \int h(\mathbf{X})\mathcal{K}_{-1}(\mathbf{X}|\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}'_3) \\ &\quad \times p(\mathbf{x}_1, \mathbf{x}_2|\mathbf{x}'_3) d\mathbf{x}_{-1} d\mathbf{x}_1 d\mathbf{x}_2, \end{aligned} \quad (4)$$

where $\tilde{\mathbb{E}}$ represents expectation with respect to $\tilde{\mathcal{K}}(\mathbf{X}|\mathbf{X}')$, $\mathcal{K}_{-1}(\mathbf{X}|\mathbf{X}')$ is the transition kernel implied by step 2 through step P of either sampler, and $\mathbf{x}_{-1} = (\mathcal{X}_{\mathcal{J}_2}, \dots, \mathcal{X}_{\mathcal{J}_P})$ is the set of components updated in step 2 through step P , which may include multiple copies of certain components of \mathbf{X} . Now the right side of (4) can be written as

$$\begin{aligned} &\tilde{\mathbb{E}}\{h(\mathbf{X})|\mathbf{x}'_3\} \\ &= \int \left\{ \int h(\mathbf{X})\mathcal{K}_{-1}(\mathbf{X}|\mathbf{x}_1, \mathbf{x}'_2, \mathbf{x}'_3) p(\mathbf{x}_1|\mathbf{x}'_2, \mathbf{x}'_3) d\mathbf{x}_{-1} d\mathbf{x}_1 \right\} \\ &\quad \times p(\mathbf{x}'_2|\mathbf{x}'_3) d\mathbf{x}'_2 \\ &= E_\pi [\mathbb{E}\{h(\mathbf{X})|\mathbf{x}'_2, \mathbf{x}'_3\}|\mathbf{x}'_3], \end{aligned}$$

where the inner expectation is with respect to $\mathcal{K}(\mathbf{X}|\mathbf{X}')$. Thus

$$\begin{aligned} E_\pi([\tilde{\mathbb{E}}\{h(\mathbf{X})|\mathbf{x}'_3\}]^2) &= E_\pi\{(E_\pi[\mathbb{E}\{h(\mathbf{X})|\mathbf{x}'_2, \mathbf{x}'_3\}|\mathbf{x}'_3])^2\} \\ &\leq E_\pi\{E_\pi([\mathbb{E}\{h(\mathbf{X})|\mathbf{x}'_2, \mathbf{x}'_3\}]^2|\mathbf{x}'_3)\} \\ &= E_\pi([\mathbb{E}\{h(\mathbf{X})|\mathbf{x}'_2, \mathbf{x}'_3\}]^2). \end{aligned}$$

But because $\text{var}_\pi[h(\mathbf{X})]$ is the same for both kernels, the maximal autocorrelation induced by $\tilde{\mathcal{K}}(\mathbf{X}|\mathbf{X}')$ is bounded above by that of $\mathcal{K}(\mathbf{X}|\mathbf{X}')$.

This shows that the computational advantages can be achieved by successively marginalizing over the components of \mathbf{X} in any single step of a Gibbs sampler. Thus repeatedly applying Theorem 2 provides the theoretical basis for the improved convergence of PCG samplers.

5. CONCLUDING REMARKS

In this article we have presented efficient Gibbs sampling techniques developed by generalizing the composition of the conditional distributions in Gibbs samplers. Unlike ordinary Gibbs samplers, PCG samplers use incompatible conditional distributions to improve the convergence characteristics of the samplers. This generalization comes at a price; the conditional distributions that compose a PCG sampler must be performed in a certain order to maintain the target stationary distribution. We have introduced three basic prescriptive tools—marginalization, permutation, and trimming—and showed how applying sequentially these tools can transform a Gibbs sampler into a PCG sampler. The resulting PCG sampler may comprise a set of incompatible conditional distributions and generally exhibits superior convergence characteristics. This strategy was

illustrated in the fitting of three models stemming from our applied work in astronomy and multiple imputation (Park and van Dyk 2008).

[Received July 2007. Revised January 2008.]

REFERENCES

- Geman, S., and Geman, D. (1984), "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6, 721–741.
- Liu, J. S. (2001), *Monte Carlo Strategies in Scientific Computing*, New York: Springer-Verlag.
- Liu, J. S., Wong, W. H., and Kong, A. (1994), Covariance Structure of the Gibbs Sampler, With Applications to Comparisons of Estimators and Augmentation Schemes," *Biometrika*, 81, 27–40.
- Park, T., and van Dyk, D. A. (2008), "Partially Collapsed Gibbs Samplers: Illustrations and Applications," technical report, University of California, Irvine.
- van Dyk, D. A., and Meng, X.-L. (1997), "Some Findings on the Orderings and Groupings of Conditional Maximizations Within ECM-Type Algorithms," *The Journal of Computational and Graphical Statistics*, 6, 202–223.