CUTTING IN LINE: SCHEDULER-CORRECTED GIBBS SAMPLING FOR STATE-DEPENDENT SCAN ORDERS

Anonymous authors
Paper under double-blind review

ABSTRACT

The Chinese Restaurant Process (CRP) provides a powerful prior for clustering in applications such as topic modeling, entity resolution, and cell-type discovery, but its standard sequential inference relies on exchangeability (symmetry under permutations) to remain valid. In modern large-scale settings, inference must be distributed and parallel. In such architectures, model state and data are often constrained by physical placement, enforcing strict locality and inducing a state-dependent traversal. We view this traversal rule as a scan scheduler: a mechanism that selects the order in which latent assignments are updated. The locality-driven scheduler studied here is the Same-Group-First (SGF) scan order. We show that SGF causes conventional parallel samplers to systematically underestimate the number of clusters. This bias worsens with dataset size, making common approximations increasingly unreliable and creating a fundamental barrier to scalable CRP inference. We then provide a theoretical explanation: SGF breaks detailed balance via two complementary effects—(i) a microscopic supermartingale drift that accelerates small-cluster depletion, and (ii) a macroscopic combinatorial fragmentation diagnostic that suppresses transitions toward more fragmented partitions. Motivated by this analysis, we propose Scheduler-Corrected Gibbs Sampling (SCG), a full-sweep Gibbs framework that treats the scan order as an auxiliary proposal variable and corrects scheduler state-dependence through an exact reverse-scan Metropolis–Hastings ratio. We instantiate SCG with a Soft-SGF Plackett-Luce scheduler, which interpolates between random scan and SGF-like same-cluster locality while preserving full reverse support. The resulting scheduler-corrected Gibbs procedure restores detailed balance while retaining a tunable preference for locality. With the empirical scaling rule $\lambda_N = c/\sqrt{N}$, SCG stabilizes this locality-reversibility tradeoff as the number of customers grows.

1 INTRODUCTION

The Chinese Restaurant Process (CRP) is a standard prior for clustering in Bayesian nonparametric models because it allows the number of clusters to grow with the data rather than being fixed in advance. A flexible prior such as the CRP allows the model to adapt as new structure emerges, a property that is essential in settings such as deep generative modeling (Falck et al., 2021), dynamic networks (Ludkin et al., 2018), federated learning (Yurochkin et al., 2019), and single-cell genomics (Borgsmüller et al., 2020). This flexibility is a direct consequence of exchangeability: the probability of a partition depends only on how data are grouped, not on the order in which observations are processed (Aldous, 2006; Pitman, 1995; 2006). In practice, inference under the CRP is almost always carried out using Gibbs sampling (Neal, 2000), where each data point is reassigned conditional on the current configuration of all others. The correctness of this procedure relies on a simple but crucial condition: the scan rule that selects the update order must not depend on the current partition in a way that changes the transition kernel. Classical implementations satisfy this by fixing a permutation in advance or by drawing one at random.

While the CRP is conceptually clean, scaling its inference to massive datasets quickly runs into a practical wall. In many nonparametric models—such as Hierarchical Dirichlet Processes (HDP) (Teh et al., 2006) or Mixed-Membership Stochastic Blockmodels (MMSB) (Airoldi et al., 2008)—each cluster carries its own parameter vector θ_k, and these parameters can be extremely high-dimensional. In large-scale topic models, for example, θ_k is a distribution over a vocabulary that can easily exceed

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

10^6 tokens (Falck et al., 2021). Once the data are distributed, synchronizing these θ_k across workers can dominate the local cost of Gibbs updates (Zinkevich et al., 2010; Ahn et al., 2015). This is why modern probabilistic systems treat distributed inference as a necessity rather than a convenience—the key challenge is no longer "how to sample," but "how to sample without paying a prohibitive communication bill" (Paritosh et al., 2025; Campbell et al., 2015).

Because synchronizing large cluster parameters across machines is so expensive, modern distributed systems are forced to organize inference around data locality rather than random access (Hennessy & Patterson, 2011; Drepper, 2007; Low et al., 2012; Liu & Wright, 2015). In practice, this means that all data points currently assigned to a cluster k, together with the cluster parameters θ_k, are placed on the same worker. Each processor therefore only needs to access the parameters stored in its own memory, greatly reducing cross-worker synchronization. It arises whenever large-scale systems organize computation around co-located groups of strongly coupled variables, as in parallel Gibbs samplers for Dirichlet-process mixtures (Williamson et al., 2013) and graph engines such as PowerGraph (Gonzalez et al., 2012; Li et al., 2023). Once the data are laid out this way, the execution order is effectively determined by memory: the system naturally updates all points belonging to one cluster before moving on to the next. We refer to the rule that selects such an update order as a *scan scheduler*. Formally, a scheduler is a distribution $s(\sigma \mid \mathbf{z})$ over permutations σ of the customer indices. The locality-driven scheduler induced by this memory layout is the Same-Group-First (SGF) scan order.

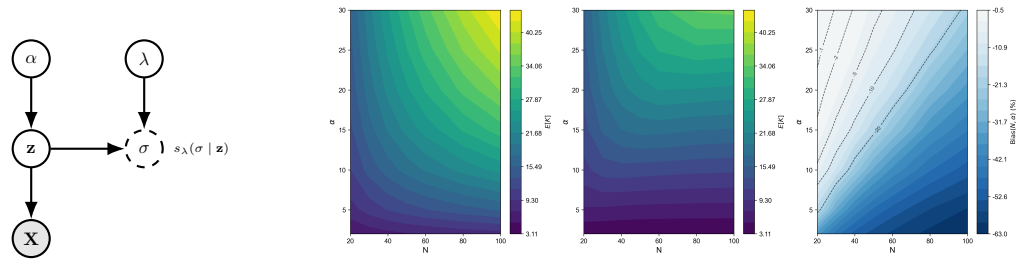
This is where a subtle conflict appears. Standard Gibbs sampling is only guaranteed to target the correct posterior when the scan scheduler is external to the current state $\mathbf{z}$ (Robert et al., 2004). Classical implementations satisfy this with state-independent schedulers, either by fixing a permutation in advance or by drawing one at random. SGF, however, is generated by the partition itself: at sweep m, the schedule σ_m is a deterministic function of the current partition $\mathbf{z}_m$. When a data point changes its cluster, it is physically moved to a different worker, and the future update order changes with it. The algorithmic path is therefore no longer external to the model—it is coupled to the evolving partition. Once the update schedule becomes part of the model state, the dynamics of the Markov chain are no longer governed by a fixed exchangeable posterior. Instead, each step reshapes the geometry that controls the next one. The resulting Markov chain is therefore no longer guaranteed to preserve the CRP posterior (Roberts & Rosenthal, 2015; He et al., 2016; De Sa et al., 2016; Terenin et al., 2020)—and, crucially, this distortion does not disappear as more data are added. It becomes stronger with scale, leading to an increasingly severe underestimation of model complexity. This makes common approximations increasingly unreliable and creates a fundamental barrier to scalable CRP inference.

We trace this behavior to a breakdown of detailed balance under state-dependent scheduling, driven by two complementary effects: (i) a supermartingale drift in cluster size induced by contiguous within-cluster updates—once a customer leaves, subsequent customers in the same block become less likely to stay, accelerating the extinction of small clusters; and (ii) a macroscopic combinatorial fragmentation diagnostic showing that locality-compatible layouts become increasingly sparse for more fragmented partitions. Together, these mechanisms show that locality changes the transition kernel itself. This motivates treating the scheduler as part of the proposal rather than as an external implementation detail.

We propose the **S**cheduler-**C**orrected **G**ibbs Sampling (**SCG**), a full-sweep Gibbs framework for state-dependent scan schedulers. SCG first draws a scan order $\sigma \sim s(\cdot \mid \mathbf{z})$, performs a complete single-site Gibbs sweep along σ, and then corrects the proposal by comparing the reverse scan probability under the proposed state with the forward scan probability under the current state. The resulting acceptance probability is

$$\min\left\{1, \frac{s(\sigma^{\leftarrow} \mid \mathbf{z}')}{s(\sigma \mid \mathbf{z})}\right\},$$

where $\sigma^{\leftarrow}$ denotes the reversed order. This ratio is exact because the single-site conditional probabilities telescope with the target-density ratio. To retain the locality motivation of SGF while avoiding its reverse-support collapse, we instantiate SCG with a Soft-SGF Plackett–Luce scheduler (Luce et al., 1959; Plackett, 1975). Viewed as an informed local proposal in a discrete state space, Soft-SGF favors consecutive updates of customers currently assigned to the same cluster while assigning positive probability to every permutation (Zanella, 2020). Its scan concentration parameter controls



(a) Soft-SGF scan scheduler.

(b) Visualization of global bias under SGF.

Figure 1: Overview of the state-dependent scan setting and its induced bias. (a) Graphical representation of the Soft-SGF scan scheduler. The CRP prior generates the partition $\mathbf{z}$ from α, while the auxiliary scan order σ is drawn from $s_\lambda(\sigma \mid \mathbf{z})$. The observed data $\mathbf{X}$ are shaded. (b) Visualization of global bias. Left: theoretical $\mathbb{E}_{\mathrm{CRP}}[K_N]$ under the CRP. Center: empirical $\widehat{\mathbb{E}}_{\mathrm{SGF}}[K_N]$ under SGF. Right: SGF relative bias $\mathrm{Bias}(N, \alpha) = 100\% \cdot (\widehat{\mathbb{E}}_{\mathrm{SGF}}[K_N] - \mathbb{E}_{\mathrm{CRP}}[K_N])/\mathbb{E}_{\mathrm{CRP}}[K_N]$.

the tradeoff between locality and reversibility: at zero it recovers random-scan Gibbs, while larger values increasingly approximate SGF.

This work makes three contributions. First, we identify state-dependent scan scheduling as a source of distortion in CRP inference and show that persistent SGF can produce severe under-fragmentation that worsens with scale. Second, we explain this behavior through accumulated within-block occupancy drift and a global combinatorial fragmentation diagnostic for locality-compatible layouts. Third, we propose Scheduler-Corrected Gibbs Sampling (SCG), an exact inference framework that treats the scan order as an auxiliary proposal variable and uses a reverse-scan Metropolis–Hastings correction. With a Soft-SGF Plackett–Luce scheduler, SCG provides controllable same-cluster locality; our experiments characterize the resulting bias–reversibility tradeoff, verify scheduler-induced locality directly, and show that an empirical rule $\lambda_N = c/\sqrt{N}$ avoids near-zero acceptance over the tested scales.

2 Theoretical Analysis and Empirical Verification

2.1 State-Dependent Scanning Violates Detailed Balance

We consider a Chinese Restaurant Process (CRP) with concentration parameter $\alpha > 0$ and N customers. Let z denote a cluster-assignment vector and let $\pi(z) = p_{\mathrm{CRP}}(z)$ denote the target exchangeable partition probability function. For a scan order $\sigma = (i_1, \ldots, i_N)$, denote by $\mathcal{K}_\sigma$ the one-sweep Gibbs transition kernel induced by updating customers in the order σ. A scan scheduler is a rule for selecting σ, represented as a distribution $s(\sigma \mid \mathbf{z})$ over permutations of $[N]$.

When the scheduler is independent of the current state—for example, when a permutation is fixed in advance or drawn independently of $\mathbf{z}$—the resulting Gibbs chain preserves p_{CRP} (Robert et al., 2004). A fixed scan need not be reversible, but it is invariant under the usual Gibbs construction. If the scan order is resampled uniformly from all permutations at each sweep, the averaged kernel is reversible because every scan order is paired with its reverse order with the same probability.

The SGF regime differs in exactly the way that matters: the scheduler is generated by the current partition. Under strict data locality, the system updates customers cluster-by-cluster, so the scan order becomes a deterministic function $\sigma_{\mathrm{SGF}}(\mathbf{z})$. Equivalently, SGF uses the state-dependent scheduler

$$s_{\mathrm{SGF}}(\sigma \mid \mathbf{z}) = \mathbf{1}\{\sigma = \sigma_{\mathrm{SGF}}(\mathbf{z})\}.$$

Specifically, customers assigned to the same persistent table are processed contiguously, with all customers in one table updated before moving to the next table. Table identities persist across updates, so changes in cluster assignments alter the SGF schedule used in subsequent sweeps. Consequently, the forward transition from $\mathbf{z}$ uses $\mathcal{K}_{\sigma_{\mathrm{SGF}}(\mathbf{z})}$, while the reverse transition from $\mathbf{z}'$ uses $\mathcal{K}_{\sigma_{\mathrm{SGF}}(\mathbf{z}')}$. Since these two scan orders are generally not reverse pairs, detailed balance with respect to p_{CRP} is no longer guaranteed. We compare SGF to a state-independent sequential Gibbs sampler and quantify bias relative to the CRP expectation. Let $K_N = K(\mathbf{z})$ denote the number of occupied clus-

ters among N customers. Figure 1(b) shows that SGF systematically underestimates the number of occupied clusters, and the bias becomes more pronounced as N increases.

In the following subsections, we analyze two complementary mechanisms—supermartingale drift and combinatorial fragmentation diagnostic—that drive this underestimation.

2.2 Supermartingale Drift and Cluster Depletion

We now formalize the microscopic mechanism behind SGF: contiguous within-cluster updates allow a negative drift in cluster size to accumulate over consecutive updates, which in turn accelerates the extinction of small clusters. We analyze a single SGF sweep in which the scan order is fixed at the start of the sweep by the current partition.

Fix a table k whose SGF block contains n_k customers at the beginning of the sweep. These n_k customers are updated consecutively. Let $\{M_j\}_{j=0}^{n_k}$ denote the occupancy of table k after j updates within this block, where M_0 is the occupancy of table k immediately before the block is processed. In general, M_0 may exceed n_k, because customers processed in earlier SGF blocks may have been reassigned to table k before its block is reached. At step $j+1$, the next customer in the run is removed from cluster k and then resampled from the CRP conditional. It returns to cluster k with probability proportional to the number of remaining customers in that cluster:

$$M_{j+1} = M_j - 1 + \mathbb{I}_{j+1}, \qquad \mathbb{I}_{j+1} \sim \text{Bernoulli}\left(\frac{M_j - 1}{\alpha + N - 1}\right), \tag{1}$$

where $\mathbb{I}_{j+1} = 1$ indicates that the updated customer returns to the original cluster k.

Lemma 1. *Within the SGF block of table k, let $\mathcal{F}_j$ be the filtration generated by the first j updates. Then $\{M_j\}_{j=0}^{n_k}$ is a supermartingale with uniformly negative conditional drift. In particular,*

$$\mu_j := \mathbb{E}[M_{j+1} - M_j \mid \mathcal{F}_j] \le -\delta, \qquad \delta := \frac{\alpha}{\alpha + N - 1}.$$

Theorem 1. *For any $1 \le L \le n_k$ and any $u > 0$,*

$$\Pr(M_L - M_0 > -L\delta + u) \le \exp\left(-\frac{2u^2}{L}\right),$$

where $\delta = \alpha/(\alpha + N - 1)$.

The proofs of Lemma 1 and Theorem 1 are provided in Appendix A.1. Lemma 1 shows that SGF imposes a negative drift on the size of the cluster being updated. Theorem 1 shows that this drift accumulates over consecutive SGF updates; the complete-run specialization is given in Appendix A.1. The result does not claim that every small cluster disappears in one sweep. It identifies the local mechanism by which SGF favors cluster loss: if early customers leave cluster k, then later customers from the same cluster are updated against a smaller remaining count, making them less likely to return. Repeated across clusters and sweeps, this depletion effect provides a local mechanism consistent with the underestimation of $\mathbb{E}[K_N]$ observed in Section 2.1.

2.3 Fragmentation Diagnostic

The drift mechanism in Section 2.2 is local to a single SGF run. We next describe a global combinatorial diagnostic for the same phenomenon. The connection to SGF comes from its locality structure: SGF arranges customers from the same cluster into contiguous blocks in the scan order. By contrast, the CRP target is exchangeable and does not privilege any particular ordering of the customers, so cluster members may be arbitrarily interleaved under a state-independent ordering. Accordingly, the fragmentation ratio measures how many K-cluster partitions remain contiguous under a fixed ordering and therefore quantifies how sparse locality-compatible layouts are within the partition space. A small ratio means that most K-cluster partitions are highly fragmented relative to the locality layout.

Fix an ordering of the N customers. Let $\mathcal{P}_{N,K}$ denote the set of all partitions of N items into K nonempty subsets, with $|\mathcal{P}_{N,K}| = S(N,K)$, the Stirling number of the second kind. Let $\mathcal{C}_{N,K} \subset \mathcal{P}_{N,K}$ denote the subset of contiguous partitions with respect to this ordering, namely partitions whose K blocks are intervals. A contiguous partition is specified by choosing $K-1$ cut points between

consecutive indices, hence $|\mathcal{C}_{N,K}| = \binom{N-1}{K-1}$. We define the fragmentation ratio

$$\rho_N(K) := \frac{|\mathcal{C}_{N,K}|}{|\mathcal{P}_{N,K}|} = \frac{\binom{N-1}{K-1}}{S(N,K)}.$$

This ratio quantifies the combinatorial compatibility between a K-cluster partition and a contiguous locality layout. A small value indicates that locality-compatible realizations occupy only a tiny fraction of the partition space.

Proposition 1. *Fix $K \geq 2$. For fixed K, as $N \to \infty$, the fragmentation ratio $\rho_N(K)$ decays exponentially in N:*

$$\rho_N(K) = \Theta_K\left(K^{-N} N^{K-1}\right).$$

In particular, $\log \rho_N(K) = -N \log K + O(\log N)$.

A proof based on the fixed-K asymptotics of the Stirling numbers is given in Appendix A.2. Proposition 1 shows that locality-compatible partitions become exponentially rare as N grows, revealing an increasingly severe combinatorial mismatch between contiguous locality layouts and the full partition space. To quantify the strength of this combinatorial bottleneck, we summarize the fragmentation ratio through the expected fragmentation severity $\Phi_{\mathrm{frag}}(N, \alpha) = \mathbb{E}_{K \sim p_{\mathrm{CRP}}}[-\log \rho_N(K)]$. Large $\Phi_{\mathrm{frag}}(N, \alpha)$ indicates stronger fragmentation pressure against locality-contiguous execution. We empirically validate this mechanism by relating $\Phi_{\mathrm{frag}}(N, \alpha)$ to the cluster retention ratio $\frac{\mathbb{E}_{\mathrm{SGF}}[K_N]}{\mathbb{E}_{\mathrm{CRP}}[K_N]}$ under controlled variations of N and α. At fixed $N = 50$ while varying α, the relationship achieves $R^2 = 0.933$; at fixed $\alpha = 10$ while varying N, it achieves $R^2 = 0.983$. These results show that increasing fragmentation severity closely tracks stronger SGF-induced cluster loss, supporting its use as a global diagnostic of the under-fragmentation effect.

3 Scheduler-Corrected Gibbs Sampling

The analysis above shows that the failure of SGF is not caused by an invalid single-site conditional update, but by the fact that the scan order itself is a function of the current partition. We therefore treat the scan order as part of the proposal. This leads to a full-sweep Metropolis–Hastings correction for state-dependent scan orders.

Let $\sigma = (i_1, \ldots, i_N)$ be a permutation of the customer indices. Given a state-dependent scan distribution $s(\sigma \mid \mathbf{z})$, a proposal from $\mathbf{z}$ is generated as follows. First draw $\sigma \sim s(\cdot \mid \mathbf{z})$. Set $\mathbf{z}^{(0)} = \mathbf{z}$, and perform a full single-site Gibbs sweep along σ: $z_{i_t}^{(t)} \sim \pi(\cdot \mid \mathbf{z}_{-i_t}^{(t-1)}), \mathbf{z}_{-i_t}^{(t)} = \mathbf{z}_{-i_t}^{(t-1)}, t = 1, \ldots, N$. The proposed state is $\mathbf{z}' = \mathbf{z}^{(N)}$. When $s(\sigma \mid \mathbf{z})$ is independent of $\mathbf{z}$, this is an ordinary Gibbs sweep. When s is state-dependent, the proposal must be corrected. Let $\sigma^{\leftarrow} = (i_N, \ldots, i_1)$ denote the reversed scan order. We accept the proposed sweep with probability

$$a(\mathbf{z}, \mathbf{z}', \sigma) = \min\left\{1, \frac{s(\sigma^{\leftarrow} \mid \mathbf{z}')}{s(\sigma \mid \mathbf{z})}\right\}. \tag{2}$$

The following result gives exactness.

Proposition 2. *Assume $s(\sigma \mid \mathbf{z}) > 0$ for all scan orders that may be used by the proposal, and that each single-site update samples exactly from $\pi(z_i \mid \mathbf{z}_{-i})$. Then the transition kernel defined by the full-sweep proposal and the acceptance probability in Eq. equation 2 satisfies detailed balance with respect to π.*

The proof follows by pairing each realized forward Gibbs path with the reversed path under $\sigma^{\leftarrow}$; the single-site conditional ratios telescope with the target-density ratio, leaving only the scheduler ratio in Eq. equation 2. The full proof is given in Appendix A.3.

3.1 Soft-SGF Plackett–Luce Scheduler

A deterministic SGF scan has poor reverse support: after a sweep changes the partition, the reverse order $\sigma^{\leftarrow}$ is generally not the SGF order under $\mathbf{z}'$, making $s(\sigma^{\leftarrow} \mid \mathbf{z}') = 0$. We therefore use a soft SGF distribution over permutations that preserves full support while favoring same-cluster runs.

Algorithm 1 Scheduler-Corrected Gibbs with Soft-SGF Scheduler

Require: partition $\mathbf{z}$, concentration α, scan parameter λ, sweeps T
1: **for** $m = 1, \dots, T$ **do**
2: Draw $\sigma \sim s_\lambda(\cdot \mid \mathbf{z})$ using Eq. equation 3; set $\widetilde{\mathbf{z}} \leftarrow \mathbf{z}$.
3: **for** i in σ **do**
4: Sample $\widetilde{z}_i \sim \pi(\cdot \mid \widetilde{\mathbf{z}}_{-i})$.
5: **end for**
6: Set $\mathbf{z}' \leftarrow \widetilde{\mathbf{z}}$ and $\ell_{\mathrm{MH}} \leftarrow \log s_\lambda(\sigma^{\leftarrow} \mid \mathbf{z}') - \log s_\lambda(\sigma \mid \mathbf{z})$.
7: With probability $\min\{1, e^{\ell_{\mathrm{MH}}}\}$, set $\mathbf{z} \leftarrow \mathbf{z}'$.
8: **end for**
9: **return** samples of $\mathbf{z}$

Given $\lambda \geq 0$, the first customer is selected uniformly from $[N] = \{1, \dots, N\}$. For $t \geq 2$, let $R_t = [N] \setminus \{i_1, \dots, i_{t-1}\}$ be the set of remaining customer indices. Conditional on the previous customer i_{t-1} and the history $i_{<t} = (i_1, \dots, i_{t-1})$, choose $i_t = j \in R_t$ with probability

$$\Pr(i_t = j \mid i_{<t}, \mathbf{z}) = \frac{\exp\{\lambda \mathbf{1}(z_j = z_{i_{t-1}})\}}{\sum_{\ell \in R_t} \exp\{\lambda \mathbf{1}(z_\ell = z_{i_{t-1}})\}}. \tag{3}$$

Here $\mathbf{1}(\cdot)$ denotes the indicator function. Thus

$$s_\lambda(\sigma \mid \mathbf{z}) = \frac{1}{N} \prod_{t=2}^{N} \frac{\exp\{\lambda \mathbf{1}(z_{i_t} = z_{i_{t-1}})\}}{\sum_{\ell \in R_t} \exp\{\lambda \mathbf{1}(z_\ell = z_{i_{t-1}})\}}.$$

This is a run-biased Plackett–Luce scan distribution. When $\lambda = 0$, all remaining customers have equal weight and $s_0(\sigma \mid \mathbf{z}) = 1/N!$, so the method reduces to an ordinary random-scan Gibbs sweep with acceptance probability one. As λ increases, the scan increasingly favors consecutive updates of customers currently assigned to the same cluster, approaching the locality pattern of SGF while retaining positive probability for every permutation.

For the CRP prior experiments, the single-site conditional in Algorithm 1 is the usual CRP reassignment conditional. After removing customer i, assigning it to an existing cluster k has weight n_k^{-i}, while creating a new cluster has weight α. In posterior models, these weights are multiplied by the corresponding predictive likelihood. The MH correction does not change: it depends only on the scan distribution s_λ, because the exact single-site conditional terms cancel in the ratio above.

3.2 Scan-Augmented Target and Marginals

The scan order is not part of the CRP data-generating model; it is an auxiliary variable introduced by the sampler. To make this distinction explicit, we consider the scan-augmented target $\widetilde{\pi}_\lambda(\mathbf{z}, \sigma) = \pi(\mathbf{z}) s_\lambda(\sigma \mid \mathbf{z}), \sigma \in \mathfrak{S}_N$, where $\mathfrak{S}_N$ is the permutation group of the N customer indices and $s_\lambda(\cdot \mid \mathbf{z})$ is the Soft-SGF scheduler. For each fixed partition $\mathbf{z}$, $s_\lambda(\cdot \mid \mathbf{z})$ is normalized over $\mathfrak{S}_N$. Therefore,

$$\sum_{\sigma \in \mathfrak{S}_N} \widetilde{\pi}_\lambda(\mathbf{z}, \sigma) = \sum_{\sigma \in \mathfrak{S}_N} \pi(\mathbf{z}) s_\lambda(\sigma \mid \mathbf{z}) = \pi(\mathbf{z}) \sum_{\sigma \in \mathfrak{S}_N} s_\lambda(\sigma \mid \mathbf{z}) = \pi(\mathbf{z}).$$

Thus the augmentation is marginally exact in $\mathbf{z}$: adding the auxiliary scan order does not alter the CRP target over partitions. This marginal identity uses only the normalization of $s_\lambda(\cdot \mid \mathbf{z})$. It should not be confused with invariance of the Markov chain, which is established by the reverse-scan Metropolis–Hastings correction in Proposition 2. The marginal distribution of the scan order is

$$p_\lambda(\sigma) = \sum_{\mathbf{z}} \widetilde{\pi}_\lambda(\mathbf{z}, \sigma) = \sum_{\mathbf{z}} \pi(\mathbf{z}) s_\lambda(\sigma \mid \mathbf{z}) = \mathbb{E}_{\mathbf{z} \sim \pi}\left[s_\lambda(\sigma \mid \mathbf{z})\right].$$

Thus the scan-order marginal is the target average of the state-dependent scan probability. Under the CRP prior, the scan-order marginal is uniform because the CRP is exchangeable and the Soft-SGF scheduler is permutation-equivariant. Let τ be any relabeling of the customer indices. Exchangeability gives $\pi(\tau \mathbf{z}) = \pi(\mathbf{z})$. Permutation equivariance of the Soft-SGF scheduler gives $s_\lambda(\tau\sigma \mid \tau \mathbf{z}) = s_\lambda(\sigma \mid \mathbf{z})$, where $\tau\sigma = (\tau(i_1), \dots, \tau(i_N))$ for $\sigma = (i_1, \dots, i_N)$. For any two scan

orders $\sigma, \sigma' \in \mathfrak{S}_N$, there exists a relabeling τ such that $\sigma' = \tau\sigma$. Using the change of variables $\mathbf{z} \mapsto \tau\mathbf{z}$, we obtain

$$p_\lambda(\sigma') = \sum_{\mathbf{z}} \pi(\mathbf{z}) s_\lambda(\tau\sigma \mid \mathbf{z}) = \sum_{\mathbf{z}} \pi(\tau\mathbf{z}) s_\lambda(\tau\sigma \mid \tau\mathbf{z}) = \sum_{\mathbf{z}} \pi(\mathbf{z}) s_\lambda(\sigma \mid \mathbf{z}) = p_\lambda(\sigma).$$

Hence all scan orders have the same marginal probability. Since $|\mathfrak{S}_N| = N!$, $p_\lambda(\sigma) = \frac{1}{N!}, \sigma \in \mathfrak{S}_N$. The uniform marginal does not make Soft-SGF equivalent to random scan. When $\lambda = 0$, $s_0(\sigma \mid \mathbf{z}) = \frac{1}{N!}$, so σ is independent of $\mathbf{z}$. When $\lambda > 0$, however, $s_\lambda(\sigma \mid \mathbf{z})$ depends on the current partition: scan orders with longer same-cluster runs receive larger conditional probability under partitions that support such runs. Thus, under the CRP prior, Soft-SGF has a uniform scan-order marginal, but σ and $\mathbf{z}$ are generally dependent.

In posterior models with observed data $\mathbf{X}$, the target becomes $\pi(\mathbf{z}) = p(\mathbf{z} \mid \mathbf{X}) \propto p_{\mathrm{CRP}}(\mathbf{z}) p(\mathbf{X} \mid \mathbf{z})$. For a fixed observed dataset $\mathbf{X}$, the likelihood generally breaks exchangeability with respect to arbitrary customer relabelings. The scan-order marginal is then

$$p_\lambda(\sigma \mid \mathbf{X}) = \sum_{\mathbf{z}} p(\mathbf{z} \mid \mathbf{X}) s_\lambda(\sigma \mid \mathbf{z}) = \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{X})} \left[s_\lambda(\sigma \mid \mathbf{z}) \right],$$

which need not be uniform. In both the prior and posterior settings, the reverse-scan Metropolis–Hastings correction accounts for the state-dependent conditional scan probability through the ratio $\frac{s_\lambda(\sigma^{\leftarrow} | \mathbf{z}')}{s_\lambda(\sigma | \mathbf{z})}$.

4 Experiments

We evaluate SCG on controlled CRP experiments and a posterior clustering task. The CRP setting isolates scan-order effects, provides the closed-form reference $\mathbb{E}_{\mathrm{CRP}}[K_N]$, and enables finite-state exactness checks. We examine persistent-SGF under-fragmentation, separate scheduler softening from MH correction through an uncorrected Soft-SGF ablation, and characterize the bias–reversibility tradeoff under fixed and scaled Soft-SGF concentrations. We then verify scheduler-induced locality directly and test whether the same behavior persists in posterior clustering.

Experimental setup. All synthetic experiments use the prior-only target $p_{\mathrm{CRP}}(\mathbf{z})$. Unless otherwise stated, fixed-λ runs use $N = 50$, $\alpha = 10$, 3000 sweeps, 500 burn-in sweeps, and 20 chains, with $\lambda \in \{0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 2.0\}$. The scaling experiment uses $\alpha = 10$, $N \in \{50, 100, 200, 500\}$, and $\lambda_N = c/\sqrt{N}$ for $c \in \{2, 3, 4\}$. Run budgets decrease with N: $2000/500/10$ sweeps/burn-in/chains for $N \leq 100$, $1500/500/8$ for $N = 200$, and $1000/300/5$ for $N = 500$. The adjacency-lift diagnostics are evaluated separately on CRP partitions using the same fixed and scaled scheduler settings.

Finite-state exactness validation. We first verify exactness on finite CRP state spaces where all partitions can be enumerated. We enumerate all canonical partitions for $N = 4$ and $N = 5$, giving 15 and 52 partition states, respectively, and set $\alpha = 2$. For each setting, we construct the full transition matrix P of SCG with the Soft-SGF scheduler and compare it against the CRP target distribution π. We report three numerical diagnostics: the maximum row-sum error of P, the stationarity error $\|\pi P - \pi\|_\infty$, and the maximum detailed-balance error $\max_{\mathbf{z},\mathbf{z}'} |\pi(\mathbf{z})P(\mathbf{z}, \mathbf{z}') - \pi(\mathbf{z}')P(\mathbf{z}', \mathbf{z})|$. Table 1 shows that all errors are at machine precision, confirming that the reverse-scan correction is exact rather than heuristic.

SGF under-fragmentation and correction. We next examine the main failure mode of uncorrected persistent SGF. Figure 2 plots the post-burn-in trajectory of K_m, the number of occupied clusters after sweep m, for $N = 50$ and $\alpha = 10$. The dashed horizontal line denotes the CRP reference $\mathbb{E}_{\mathrm{CRP}}[K_N]$, and the gray curve shows the persistent SGF baseline. Persistent SGF remains far below the CRP reference, confirming severe under-fragmentation under persistent SGF. SCG corrects this distortion when the Soft-SGF scheduler is not too concentrated: for $\lambda = 0.5$ and 1.0, the traces remain close to the CRP reference, whereas at $\lambda = 2.0$ the reverse-scan acceptance probability nearly degenerates and the chain becomes trapped at very small K_m.

Meanwhile, we also evaluate the fixed-λ tradeoff more systematically. Figure 3a–3b summarizes this behavior: the relative bias remains close to zero for moderate λ, while the MH acceptance rate decreases steadily as the scheduler becomes more concentrated. At large λ, near-zero acceptance

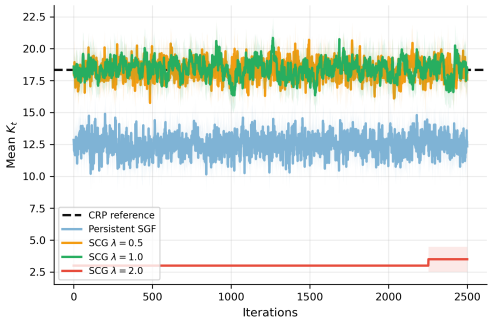


Figure 2: Post-burn-in trace of K_m under persistent SGF and Soft-SGF SCG. The dashed line denotes $\mathbb{E}_{\mathrm{CRP}}[K_N]$, and bias values are relative to the CRP reference.

Table 1: Finite-state exactness validation. All row-sum, stationarity, and detailed-balance errors are at machine precision.

N	α	λ	States	**Row err.**	**Stat. err.**	**DB err.**
4	2.0	0.0	15	$4.4\mathbf{e{-16}}$	$5.6\mathbf{e{-17}}$	$6.9\mathbf{e{-18}}$
4	2.0	0.5	15	$4.4\mathbf{e{-16}}$	$5.6\mathbf{e{-17}}$	$6.9\mathbf{e{-18}}$
4	2.0	1.0	15	$1.1\mathbf{e{-15}}$	$8.3\mathbf{e{-17}}$	$6.9\mathbf{e{-18}}$
5	2.0	0.5	52	$1.1\mathbf{e{-15}}$	$5.6\mathbf{e{-17}}$	$4.3\mathbf{e{-18}}$
5	2.0	1.0	52	$2.9\mathbf{e{-15}}$	$1.1\mathbf{e{-16}}$	$5.2\mathbf{e{-18}}$

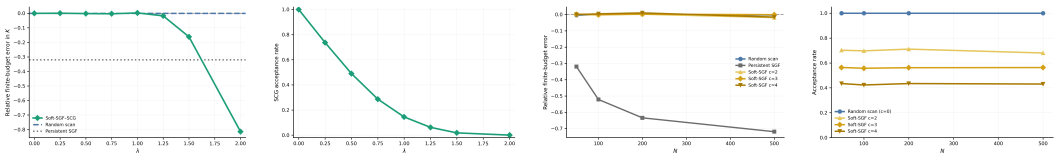
leads to poor finite-budget mixing and substantial empirical bias, although the invariant distribution remains exact by Proposition 2.

Uncorrected Soft-SGF ablation. To separate scheduler softening from the Metropolis–Hastings correction, we additionally run the same Soft-SGF scheduler without the reverse-scan acceptance step. Softening alone substantially reduces the severe distortion of persistent SGF: the uncorrected sampler has relative errors of $+0.68\%$, $+1.90\%$, and $+4.07\%$ at $\lambda = 0.5, 1.0$, and 2.0, respectively, compared with -32.1% for persistent SGF. However, because the scheduler remains state-dependent, this variant does not in general preserve the CRP target. Where acceptance remains practical, SCG reduces the observed deviations to -0.19% and $+0.27\%$ at $\lambda = 0.5$ and 1.0, respectively. At $\lambda = 2$, the SCG acceptance rate falls to 2×10^{-5}, showing that overly aggressive locality can make exact full-sweep correction impractical under a finite simulation budget.

Scaling the scheduler concentration. Motivated by the acceptance collapse at large λ, we next scale the scheduler concentration with N. Because the reverse-scan log ratio accumulates over a full sweep of N single-site updates, a fixed λ that is usable at small N can lead to near-zero acceptance at larger scales. We therefore use the empirical scaling rule $\lambda_N = \frac{c}{\sqrt{N}}$, which weakens the per-step locality preference as the sweep length increases. Figure 3c–3d evaluates this rule over $N \in \{50, 100, 200, 500\}$ with $\alpha = 10$. Persistent SGF becomes increasingly biased as N grows, whereas scaled Soft-SGF SCG keeps the relative bias close to zero across all tested sizes. The reverse-scan acceptance rate remains non-degenerate for $c \in \{2, 3, 4\}$, avoiding the near-zero-acceptance regime over the tested range.

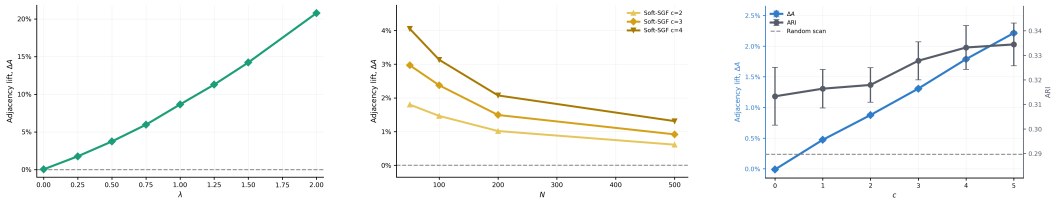
Scheduler-induced locality. We finally verify that Soft-SGF realizes the intended same-group-first preference. For a scan order σ and pre-sweep partition $\mathbf{z}$, let $A(\sigma, \mathbf{z}) = \frac{1}{N-1} \sum_{t=2}^{N} \mathbf{1}\{z_{i_t} = z_{i_{t-1}}\}$ denote the realized same-cluster adjacency. We measure scheduler-induced locality by the adjacency lift $\Delta A(\sigma, \mathbf{z}) = A(\sigma, \mathbf{z}) - A_{\mathrm{rand}}(\mathbf{z}), A_{\mathrm{rand}}(\mathbf{z}) = \frac{\sum_k n_k(n_k-1)}{N(N-1)}$, where $A_{\mathrm{rand}}(\mathbf{z})$ is the expected adjacency of a uniformly random permutation of the same pre-sweep partition. Conditioning on the same partition removes adjacency arising solely from the current cluster-size distribution. Figure 4a shows that ΔA increases monotonically with λ, from approximately zero at $\lambda = 0$ to 0.208 at $\lambda = 2$. Under $\lambda_N = c/\sqrt{N}$, Figure 4b shows positive lift for every tested $c > 0$, with larger c producing stronger locality. Thus, Soft-SGF induces controllable same-cluster locality, while the preceding results characterize the corresponding reversibility cost.

Posterior clustering on real data. We further evaluate the scan-order effect in a posterior clustering task. We use the sklearn digits dataset (Pedregosa et al., 2011), keep digits 0–9, and reduce the features to 10 principal components. We fit a spherical Dirichlet-process Gaussian mixture model with $\mathbf{z} \sim \mathrm{CRP}(\alpha)$, $\theta_k \sim \mathcal{N}(0, \kappa^2 I)$, $x_i \mid z_i = k, \theta_k \sim \mathcal{N}(\theta_k, \eta^2 I)$, using $\eta^2 = 1.0$ and $\kappa^2 = 10.0$. The cluster parameters are analytically marginalized, so each single-site update uses the collapsed posterior predictive likelihood. The preprocessed dataset contains $N = 500$ observations. For each $\alpha \in \{1, 5, 10\}$, we run 10 independent chains for 200 sweeps with no burn-in, initialized from a $K = 20$ k-means partition. For Soft-SGF SCG, we sweep $\lambda_N = c/\sqrt{N}$ over $c \in \{0, 1, 2, 3, 4, 5\}$. Figure 4(c) shows that the posterior adjacency lift increases monotonically



(a) Fixed-λ: bias (b) Fixed-λ: MH Acc. (c) Scaled λ_N: bias (d) Scaled λ_N: MH Acc.

Figure 3: Bias–reversibility tradeoffs in controlled CRP experiments. Panels (a)–(b) report relative bias and Metropolis–Hastings acceptance under fixed λ at $N = 50$ and $\alpha = 10$. Panels (c)–(d) report the corresponding quantities under the scaled rule $\lambda_N = c/\sqrt{N}$ with $\alpha = 10$.



(a) Fixed-λ: locality lift (b) Scaled λ_N: locality lift (c) Posterior locality and ARI

Figure 4: Scheduler-induced locality and posterior behavior. Panels (a)–(b) show adjacency lift under fixed and scaled Soft-SGF settings. Panel (c) shows posterior adjacency lift (left axis) and ARI (right axis) on sklearn digits under $\lambda_N = c/\sqrt{N}$.

with c, from approximately zero at $c = 0$ to 0.022 at $c = 5$, while the Adjusted Rand Index (ARI) (Hubert & Arabie, 1985) remains stable across the tested range. Thus, the controllable same-cluster locality observed in the CRP experiments persists under posterior inference. In the matched $c = 3$ comparison, persistent SGF yields mean $K_{200} = 1.20$, compared with 9.77 for random scan and 10.30 for Soft-SGF SCG. Persistent SGF therefore collapses to very few posterior clusters, whereas SCG retains a cluster count close to the random-scan reference while inducing positive scheduler-controlled locality.

5 Discussion

We identified state-dependent scan scheduling as a source of distortion in Chinese Restaurant Process inference. Persistent Same-Group-First (SGF) scanning couples the update order to the current partition and can produce severe under-fragmentation. Our analysis explains this behavior from two complementary perspectives: contiguous within-cluster updates allow negative occupancy drift to accumulate within SGF blocks, while the fragmentation diagnostic shows that locality-compatible arrangements become increasingly sparse as the partition becomes more fragmented. We proposed Scheduler-Corrected Gibbs Sampling (SCG), which treats the scan order as an auxiliary proposal variable and applies an exact reverse-scan Metropolis–Hastings correction. Finite-state enumeration verifies stationarity and detailed balance to numerical precision. Controlled CRP experiments further show that persistent SGF can substantially underestimate the number of occupied clusters, whereas SCG recovers the target behavior when reverse-scan acceptance remains practical. The uncorrected Soft-SGF ablation shows that softening reduces SGF distortion, while MH correction ensures exact target preservation. The experiments also characterize how scheduler concentration controls the tradeoff between locality and reversibility. Increasing the Soft-SGF concentration produces stronger scheduler-induced same-cluster locality, as measured by adjacency lift, but lowers reverse-scan acceptance and can impair finite-budget mixing. The empirical scaling rule $\lambda_N = c/\sqrt{N}$ avoids the near-zero-acceptance regime over the tested problem sizes while retaining positive locality lift. Posterior clustering on sklearn digits shows the same pattern: persistent SGF collapses to very few clusters, whereas Soft-SGF SCG avoids this collapse, retains controllable same-cluster locality, and maintains stable clustering quality across the tested concentration range. The current formulation focuses on full-sweep correction and a Plackett–Luce relaxation of SGF, with partial-sweep and adaptive extensions left for future work. More broadly, the framework illustrates how locality-aware execution policies can be incorporated into MCMC as proposal mechanisms rather than treated as external implementation details.

AI USE STATEMENT

Generative AI tools were used to assist with research ideation and execution, literature retrieval and discovery, drafting and polishing the manuscript, and developing and checking mathematical arguments and proofs. The authors reviewed all AI-assisted material and independently checked the final mathematical claims, cited literature, experimental results, and manuscript content. The authors take full responsibility for the final content of this work, including all text, claims, proofs, and experimental conclusions.

REPRODUCIBILITY STATEMENT

We provide the information needed to reproduce our theoretical and empirical results. Experimental settings and evaluation protocols are described in Section 4, complete proofs are given in Appendix A, and additional results and implementation details are provided in the supplementary material.

REFERENCES

Sungjin Ahn, Anoop Korattikara, Nathan Liu, Suju Rajan, and Max Welling. Large-scale distributed bayesian matrix factorization using stochastic gradient mcmc. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 9–18, 2015.

Edo M Airoldi, David Blei, Stephen Fienberg, and Eric Xing. Mixed membership stochastic blockmodels. *Advances in neural information processing systems*, 21, 2008.

David J Aldous. Exchangeability and related topics. In *École d'Été de Probabilités de Saint-Flour XIII—1983*, pp. 1–198. Springer, 2006.

Nico Borgsmüller, Jose Bonet, Francesco Marass, Abel Gonzalez-Perez, Nuria Lopez-Bigas, and Niko Beerenwinkel. Bnpc: Bayesian non-parametric clustering of single-cell mutation profiles. *Bioinformatics*, 36(19):4854–4859, 2020.

Trevor Campbell, Julian Straub, John W Fisher III, and Jonathan P How. Streaming, distributed variational inference for bayesian nonparametrics. *Advances in Neural Information Processing Systems*, 28, 2015.

Christopher De Sa, Chris Re, and Kunle Olukotun. Ensuring rapid mixing and low bias for asynchronous gibbs sampling. In *International Conference on Machine Learning*, pp. 1567–1576. PMLR, 2016.

Ulrich Drepper. What every programmer should know about memory. *Red Hat, Inc*, 11(2007):2007, 2007.

Fabian Falck, Haoting Zhang, Matthew Willetts, George Nicholson, Christopher Yau, and Chris C Holmes. Multi-facet clustering variational autoencoders. *Advances in Neural Information Processing Systems*, 34:8676–8690, 2021.

Joseph E Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. {PowerGraph}: Distributed {Graph-Parallel} computation on natural graphs. In *10th USENIX symposium on operating systems design and implementation (OSDI 12)*, pp. 17–30, 2012.

Bryan D He, Christopher M De Sa, Ioannis Mitliagkas, and Christopher Ré. Scan order in gibbs sampling: Models in which it matters and bounds on how much. *Advances in neural information processing systems*, 29, 2016.

John L Hennessy and David A Patterson. *Computer architecture: a quantitative approach*. Elsevier, 2011.

Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.

Xue Li, Ke Meng, Lu Qin, Longbin Lai, Wenyuan Yu, Zhengping Qian, Xuemin Lin, and Jingren Zhou. Flash: A framework for programming distributed graph processing algorithms. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pp. 232–244. IEEE, 2023.

Ji Liu and Stephen J Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.

Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M Hellerstein. Distributed graphlab: A framework for machine learning in the cloud. *arXiv preprint arXiv:1204.6078*, 2012.

R Duncan Luce et al. *Individual choice behavior*, volume 4. Wiley New York, 1959.

Matthew Ludkin, Idris Eckley, and Peter Neal. Dynamic stochastic block models: parameter estimation and detection of changes in community structure. *Statistics and Computing*, 28(6):1201–1213, 2018.

Radford M Neal. Markov chain sampling methods for dirichlet process mixture models. *Journal of computational and graphical statistics*, 9(2):249–265, 2000.

Parth Paritosh, Nikolay Atanasov, and Sonia Martínez. Distributed variational inference for online supervised learning. *IEEE Transactions on Control of Network Systems*, 2025.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.

Jim Pitman. Exchangeable and partially exchangeable random partitions. *Probability theory and related fields*, 102(2):145–158, 1995.

Jim Pitman. *Combinatorial stochastic processes: Ecole d'eté de probabilités de saint-flour xxxii-2002*. Springer Science & Business Media, 2006.

Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.

Christian P Robert, George Casella, and George Casella. *Monte Carlo statistical methods*, volume 2. Springer, 2004.

Gareth O Roberts and Jeffrey S Rosenthal. Surprising convergence properties of some simple gibbs samplers under various scans. *International Journal of Statistics and Probability*, 5(1):51–60, 2015.

Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006. doi: 10.1198/016214506000000302. URL https://doi.org/10.1198/016214506000000302.

Alexander Terenin, Daniel Simpson, and David Draper. Asynchronous gibbs sampling. In *International Conference on Artificial Intelligence and Statistics*, pp. 144–154. PMLR, 2020.

Sinead Williamson, Avinava Dubey, and Eric Xing. Parallel markov chain monte carlo for nonparametric mixture models. In *International Conference on Machine Learning*, pp. 98–106. PMLR, 2013.

Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *International conference on machine learning*, pp. 7252–7261. PMLR, 2019.

Giacomo Zanella. Informed proposals for local mcmc in discrete spaces. *Journal of the American Statistical Association*, 115(530):852–865, 2020.

Martin Zinkevich, Markus Weimer, Lihong Li, and Alex Smola. Parallelized stochastic gradient descent. *Advances in neural information processing systems*, 23, 2010.

A Additional Proofs

A.1 Proofs for Supermartingale Drift

Proof of Lemma 1. From Eq. equation 1,

$$M_{j+1} - M_j = -1 + \mathbb{I}_{j+1}.$$

Therefore,

$$\begin{aligned}\mu_j &= \mathbb{E}[M_{j+1} - M_j \mid \mathcal{F}_j] \\ &= -1 + \mathbb{E}[\mathbb{I}_{j+1} \mid \mathcal{F}_j] = -1 + \frac{M_j - 1}{\alpha + N - 1}.\end{aligned}$$

Since $M_j \le N$,

$$\mu_j \le -1 + \frac{N-1}{\alpha + N - 1} = -\frac{\alpha}{\alpha + N - 1} = -\delta.$$

Thus $\{M_j\}$ has non-positive conditional drift, and the displayed uniform bound proves the claim. □

Proof of Theorem 1. For $0 \le \ell \le L$, define

$$A_\ell := \sum_{j=0}^{\ell-1} \mu_j, \qquad A_0 = 0,$$

where

$$\mu_j = \mathbb{E}[M_{j+1} - M_j \mid \mathcal{F}_j].$$

Define

$$Y_\ell := M_\ell - M_0 - A_\ell.$$

Then $\{Y_\ell\}_{\ell=0}^{L}$ is a martingale, because

$$Y_{j+1} - Y_j = (M_{j+1} - M_j) - \mu_j$$

has conditional expectation zero given $\mathcal{F}_j$. Moreover,

$$M_L - M_0 = Y_L + A_L.$$

By Lemma 1, $\mu_j \le -\delta$ for all j, hence

$$A_L = \sum_{j=0}^{L-1} \mu_j \le -L\delta.$$

Let

$$\xi_j := Y_{j+1} - Y_j = (M_{j+1} - M_j) - \mu_j$$

be the martingale difference. Since $M_{j+1} - M_j \in \{-1, 0\}$, conditional on $\mathcal{F}_j$, the increment ξ_j lies in an interval of length at most one. The Hoeffding–Azuma inequality for martingales with bounded conditional ranges gives

$$\Pr(Y_L \ge u) \le \exp\left(-\frac{2u^2}{L}\right).$$

Finally, since $M_L - M_0 = Y_L + A_L$ and $A_L \le -L\delta$,

$$\{M_L - M_0 > -L\delta + u\} \subseteq \{Y_L > u\}.$$

Therefore,

$$\Pr(M_L - M_0 > -L\delta + u) \le \Pr(Y_L > u) \le \exp\left(-\frac{2u^2}{L}\right).$$

□

Corollary 1 (Accumulated drift under SGF)**.** *For any $u > 0$,*

$$\Pr(M_{n_k} > M_0 - n_k\delta + u) \le \exp\left(-\frac{2u^2}{n_k}\right).$$

Proof. Apply Theorem 1 with $L = n_k$. Then

$$M_{n_k} - M_0 > -n_k\delta + u$$

is equivalent to

$$M_{n_k} > M_0 - n_k\delta + u.$$

The result follows immediately. □

A.2 Proof of the Fragmentation Ratio Bound

Proof of Proposition 1. For fixed K, the Stirling number of the second kind satisfies

$$S(N,K) = \frac{1}{K!}\sum_{j=0}^{K}(-1)^{K-j}\binom{K}{j}j^N \sim \frac{K^N}{K!} \qquad (N \to \infty).$$

Also,

$$\binom{N-1}{K-1} \sim \frac{N^{K-1}}{(K-1)!}.$$

Therefore,

$$\rho_N(K) = \frac{\binom{N-1}{K-1}}{S(N,K)} \sim \frac{N^{K-1}}{(K-1)!}\frac{K!}{K^N} = K\frac{N^{K-1}}{K^N}.$$

Thus

$$\rho_N(K) = \Theta\big(K^{-N}N^{K-1}\big).$$

Taking logarithms gives

$$\begin{aligned}\log \rho_N(K) &= -N\log K + (K-1)\log N + O(1)\\ &= -N\log K + O(\log N).\end{aligned}$$

□

A.3 Proof of SCG Exactness

Proof of Proposition 2. The forward joint proposal probability for a realized path $\mathbf{z}^{(0)} = \mathbf{z},\ \mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N)} = \mathbf{z}'$ is

$$Q(\mathbf{z}', \sigma \mid \mathbf{z}) = s(\sigma \mid \mathbf{z})\prod_{t=1}^{N}\pi\left(z_{i_t}^{(t)} \mid \mathbf{z}_{-i_t}^{(t-1)}\right).$$

The paired reverse move starts from $\mathbf{z}'$, uses $\sigma^{\leftarrow}$, and follows the reversed path $\mathbf{z}' = \mathbf{z}^{(N)}, \mathbf{z}^{(N-1)}, \dots, \mathbf{z}^{(0)} = \mathbf{z}$. Its proposal probability is

$$Q(\mathbf{z}, \sigma^{\leftarrow} \mid \mathbf{z}') = s(\sigma^{\leftarrow} \mid \mathbf{z}')\prod_{t=1}^{N}\pi\left(z_{i_t}^{(t-1)} \mid \mathbf{z}_{-i_t}^{(t)}\right).$$

Since the t-th forward update changes only coordinate i_t,

$$\mathbf{z}_{-i_t}^{(t)} = \mathbf{z}_{-i_t}^{(t-1)}.$$

Hence

$$\frac{\pi\left(z_{i_t}^{(t-1)} \mid \mathbf{z}_{-i_t}^{(t)}\right)}{\pi\left(z_{i_t}^{(t)} \mid \mathbf{z}_{-i_t}^{(t-1)}\right)} = \frac{\pi(\mathbf{z}^{(t-1)})}{\pi(\mathbf{z}^{(t)})}.$$

Multiplying over $t = 1, \dots, N$ gives

$$\prod_{t=1}^{N} \frac{\pi\left(z_{i_t}^{(t-1)} \mid \mathbf{z}_{-i_t}^{(t)}\right)}{\pi\left(z_{i_t}^{(t)} \mid \mathbf{z}_{-i_t}^{(t-1)}\right)} = \frac{\pi(\mathbf{z})}{\pi(\mathbf{z}')}.$$

Therefore the Metropolis–Hastings ratio is

$$\frac{\pi(\mathbf{z}')Q(\mathbf{z}, \sigma^{\leftarrow} \mid \mathbf{z}')}{\pi(\mathbf{z})Q(\mathbf{z}', \sigma \mid \mathbf{z})} = \frac{s(\sigma^{\leftarrow} \mid \mathbf{z}')}{s(\sigma \mid \mathbf{z})}.$$

Thus Eq. equation 2 is exactly the Metropolis–Hastings acceptance probability for the paired forward and reverse paths, and the resulting transition kernel satisfies detailed balance with respect to π. □

B Additional Experimental Details

B.1 Posterior Clustering Traces on Digits

Figure 5 provides the sweep-by-sweep posterior cluster-count trajectories corresponding to the digits experiment in Section 4. These traces complement the aggregate results in the main text by showing the evolution of the occupied cluster count throughout posterior sampling.

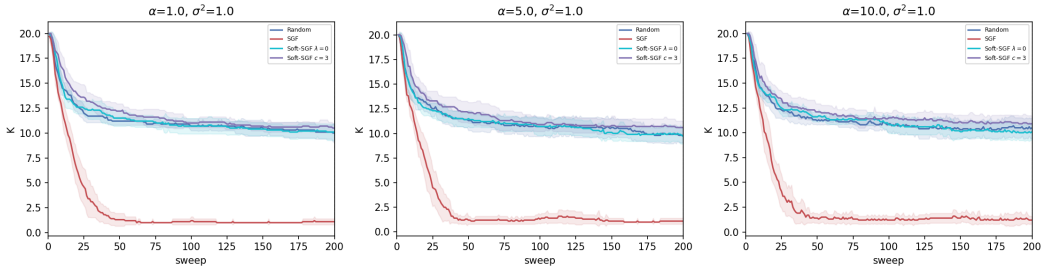


Figure 5: Posterior clustering traces on sklearn digits under a spherical DP-GMM. Panels show K_m, the number of occupied clusters after sweep m, for $\alpha \in \{1, 5, 10\}$ with $\eta^2 = 1.0$. All chains are initialized from a $K = 20$ k-means partition.

B.2 Posterior DP-GMM Implementation Details

For the posterior experiment, the Gaussian cluster means are analytically integrated out. Let $n_{k,-i}$ and

$$S_{k,-i} = \sum_{j \neq i : z_j = k} x_j$$

denote the size and sufficient statistic of cluster k after removing observation i. Under $\theta_k \sim \mathcal{N}(0, \kappa^2 I)$ and $x_i \mid z_i = k, \theta_k \sim \mathcal{N}(\theta_k, \eta^2 I)$, the posterior variance and mean of the cluster center are

$$V_{k,-i} = \left(\kappa^{-2} + n_{k,-i}\eta^{-2}\right)^{-1}, \qquad m_{k,-i} = V_{k,-i}\eta^{-2} S_{k,-i}.$$

Hence the collapsed Gibbs weight for assigning x_i to an existing cluster k is proportional to

$$n_{k,-i}\, \mathcal{N}\left(x_i; m_{k,-i}, (\eta^2 + V_{k,-i})I\right),$$

while the weight for creating a new cluster is proportional to

$$\alpha\, \mathcal{N}\left(x_i; 0, (\eta^2 + \kappa^2)I\right).$$

The same collapsed conditional is used across all scan-order variants. The methods differ only in the scan scheduler and, for SCG, the reverse-scan Metropolis–Hastings correction.