
Techniques for proving Asynchronous Convergence results for Markov Chain Monte Carlo methods

Alexander Terenin*
Petuum, Inc. and
Imperial College London
A.TERENIN17@IMPERIAL.AC.UK

Eric P. Xing
Petuum, Inc. and
Carnegie Mellon University
EPXING@CS.CMU.EDU

Abstract

Markov Chain Monte Carlo (MCMC) methods such as Gibbs sampling are finding widespread use in applied statistics and machine learning. These often require significant computational power, and are increasingly being deployed on parallel and distributed systems such as compute clusters. Recent work has proposed running iterative algorithms such as gradient descent and MCMC in parallel *asynchronously* for increased performance, with good empirical results in certain problems. Unfortunately, for MCMC this parallelization technique requires new convergence theory, as it has been explicitly demonstrated to lead to divergence on some examples. Recent theory on *Asynchronous Gibbs sampling* describes why these algorithms can fail, and provides a way to alter them to make them converge. In this article, we describe how to apply this theory in a generic setting, to understand the asynchronous behavior of any MCMC algorithm, including those implemented using parameter servers, and those not based on Gibbs sampling.

Keywords: Bayesian statistics, big data, Gibbs sampling, iterative algorithm, Metropolis-Hastings algorithm, parallel and distributed systems, parameter server.

1 Introduction

Drawing samples from a posterior probability distribution for Bayesian learning is a fundamental task in today's statistical practice, machine learning, and data science. Indeed, Bayesian methods including Gaussian Processes, Dirichlet Processes, mixed-effects regression, and many others, have found increased industrial application in recent years – Latent Dirichlet Allocation [3] is one widely-deployed example.

Unfortunately, Markov Chain Monte Carlo (MCMC) methods – the cornerstone of modern Bayesian computation over the last two decades – often do not scale well with data size or model complexity. Recent work has sought to address this issue by deploying these methods on parallel computational hardware such as GPUs [12] and compute clusters [7, 10].

In the systems community, there has been recent interest in *asynchronous* approaches to machine learning, particularly for optimization-based methods. In an asynchronous algorithm, worker nodes will perform computation *as fast as they can*, without waiting on updates from other workers to arrive. Ho et al. [6] and Wei et al. [14] have proposed a distributed system architecture called a *parameter server* for running machine learning algorithms in this fashion. Asynchronous methods can utilize hardware more effectively than fully synchronous methods, because worker nodes don't spend time waiting for other nodes, but require detailed attention because their convergence is in general not implied by standard optimization and Monte Carlo theory.

*Website: <http://avt.im/>




Figure 1: Iterations of fully synchronous MCMC and asynchronous MCMC with shared memory.

For Monte Carlo, asynchronous methods were first considered by the Latent Dirichlet Allocation (LDA) community, applied to Gibbs sampling. In particular, Newman et al. [10] proposed an *Asynchronous Gibbs sampler* for LDA [13], which was subsequently analyzed in an LDA-specific fashion by Ihler and Newman [7]. Johnson et al. [8] showed that Asynchronous Gibbs can diverge on certain Gaussian targets. The first fully generic theory developed was by Terenin et al. [11], where it was shown that if workers are allowed to occasionally reject messages from other workers in a systematic way, then Asynchronous Gibbs Sampling converges. Finally, De Sa et al. [4], whose work appeared concurrently, analyzed Asynchronous Gibbs in a completely different manner.

In this article, we outline how to apply the framework in Terenin et al. [11] beyond the context of Gibbs sampling. We prove that Asynchronous MCMC in a *shared memory* setting converges without any additional assumptions. We show how these ideas extend to the *compute cluster* setting, using the framework to prove a new convergence result for Asynchronous MCMC on parameter servers.

This article is purely theoretical: there will be no data or discussion of performance. Our aim is to present and illustrate the types of analyses that are possible with current theory.

2 Asynchronous Markov Chain Monte Carlo

Before analyzing asynchronous MCMC, we first describe notation. We begin by defining a convergent fully synchronous algorithm, on which subsequent asynchronous algorithms will be based.

Definition 1. Let Ω be a topological space representing the parameter space for the given problem. Let \mathcal{M} be the space of Borel probability measures over Ω . Let $\pi \in \mathcal{M}$ be the target probability measure, and let $\mu \in \mathcal{M}$ be an arbitrary measure. Define a Markov operator $P : \mathcal{M} \rightarrow \mathcal{M}$ satisfying

$$\|P^k(\mu) - \pi\|_{TV} \rightarrow 0 \quad (1)$$

as $k \rightarrow \infty$, where *TV* denotes total variation distance.

To unpack this definition, consider an example. If we take $\Omega = \mathbb{R}$ and π to be Gaussian, then P could be a Metropolis-Hastings algorithm targeting said Gaussian. At time $k = 0$, if μ is the distribution of the algorithm, then its distribution at time $k = 1$ is $P(\mu)$. It is important to highlight that P converges on the space of measures \mathcal{M} , not on Ω . Figure 1 illustrates one possible trajectory of P in a fully synchronous setting.

2.1 Asynchronous MCMC on multiple threads with shared memory

Now, we consider extending the above framework to the simplest asynchronous setting: a multi-threaded computer that stores the current algorithm’s state in shared memory. The process will run according to the following procedure in parallel on every thread.

Algorithm 2. For each thread do the following in parallel with no synchronization.

1. Read a value x with distribution μ from shared memory.
2. Compute the random variable x^* with distribution $\mu^* = P(\mu)$.
3. Write x^* to shared memory, so that its distribution becomes μ^* .

Algorithm 2 is illustrated in Figure 1. We may understand it from two dual perspectives.

- (a) For a fixed value $x \in \Omega$, the next value x^* is a random variable.
- (b) For a fixed measure $\mu \in \mathcal{M}$, the next measure μ^* is also fixed.




Figure 2: Iterations of an asynchronous Gibbs sampler on a compute cluster.

Hence, an MCMC algorithm is a random algorithm with respect to Ω , but a non-random algorithm with respect to \mathcal{M} . In the fully synchronous setting, perspective (a) tends to be more powerful, because it allows us to analyze P using stochastic process theory. In the asynchronous setting, working with this perspective becomes very challenging because the random variables representing states of the algorithm no longer possess the Markov property. We will thus utilize perspective (b) in this analysis. The following property will be of interest.

Result 3. *Suppose that P is a Markov operator with unique limiting distribution π and initial distribution μ . Then*

$$\|P^{k+1}(\mu) - \pi\|_{\text{TV}} \leq \|P^k(\mu) - \pi\|_{\text{TV}} \quad (2)$$

and we say that P is a contracting operator with respect to the total variation metric.

Proof. Meyn and Tweedie [9], Proposition 13.3.2. ■

For any Markov chain that converges to its stationary distribution, this implies that as $k \rightarrow \infty$, we have $\|P^k(\mu) - \pi\|_{\text{TV}} \searrow 0$. We can use this property to prove that Algorithm 2 converges.

Theorem 4. *Let P be a Markov operator with unique limiting distribution π . Let $x_0 \sim \mu_0$ for some arbitrary μ_0 . Let x_1, x_2, \dots be the sequence of values written to shared memory by the threads, and let μ_1, μ_2, \dots be their distributions. Assume that there are no partial reads or writes of x . Assume that P remains time-homogeneous if executed asynchronously. Assume that the maximum time between two updates made by the same thread is bounded above by b . Then we have that $\lim_{k \rightarrow \infty} \mu_k = \pi$.*

Proof. Appendix A. The key idea is that applying P can only reduce the distance to stationarity, which causes the asynchronous algorithm to inherit convergence from its fully synchronous counterpart. ■

Theorem 4 shows us how convergence may be analyzed in settings where we do not have the Markov property. The approach is appealing because every convergent MCMC algorithm's Markov operator is a contracting operator with respect to the total variation metric, and so it is completely general.

The main drawbacks of Theorem 4 are that it requires threads to read and write the state x in full, and that x can only exist in one central location. This limits the practical utility of Algorithm 2, because it makes it difficult to apply in a distributed setting.

2.2 Asynchronous MCMC on a compute cluster

We now show how to extend the above ideas to the compute cluster setting, where the state x may exist in multiple locations concurrently. We begin by introducing the framework of Baudet [1], Bertsekas [2], and Frommer and Szyld [5], as applied to MCMC by Terenin et al. [11].

Result 5. *Suppose that $H : E \rightarrow E$ is a weakly coupled Markov operator with unique coupled limiting distribution Π – these refer to the entire cluster, whereas π and \mathcal{M} refer to each worker. Let m be the number of workers, assumed fixed and finite. Assume the following.*

$$\text{Box Condition: we have } E = \bigtimes_{i=1}^m \mathcal{M}.$$

Let H_i be the i th component of H with respect to E – this is each worker's component of the coupled operator. Assume H is constructed from a Markov operator P such that it inherits Result 3 with respect to each H_i . Under appropriate regularity conditions given by Assumption 11 of Appendix B, H will converge asynchronously in the sense of Definition 9 of Appendix B. See also Remark 10.




Figure 3: Iterations of an asynchronous MCMC sampler on a parameter server.

Proof. Terenin et al. [11]. ■

The main challenge in using this result is satisfying the box condition, particularly since an arbitrary measure space is not a product space. One way to do this is to embed the target distribution of interest π within some larger Π that is contained in a product space, for instance by taking

$$\Pi = (\pi, \dots, \pi). \quad (3)$$

Having done this, we need to construct an H that converges to Π . We can do so using the Metropolis-Hastings method. For details, see Terenin et al. [11], where such chains are proven to converge. In that work, a convergent Asynchronous Gibbs sampler is constructed by taking H_i to be a Metropolis-Hastings proposal based on a Gibbs step – this is illustrated in Figure 2.

There are many possible alternatives. We can both choose E to accommodate different kinds of parallel and distributed systems, and H to accommodate other MCMC schemes. For instance, consider the parameter server distributed system architecture of Ho et al. [6] and Wei et al. [14]. Here, we have a main node called the parameter server that allows worker nodes to read a state x , and send updated states x^* . Upon receiving x^* , the parameter server updates its state using a pre-defined rule. Consider now the following algorithm executed on a parameter server.

Algorithm 6. For a set of workers, given an unnormalized target posterior distribution π , do the following in parallel with no synchronization.

1. Read $(x, \pi(x))$ from the parameter server, and let μ be the distribution of x .
2. Compute the random variable x^* using a valid MCMC step and evaluate $\pi(x^*)$.
3. Send $(x, x^*, \pi(x^*))$ to the parameter server. The parameter server accepts the message and sets $x_s = x^*$ with Metropolis-Hastings probability

$$\min \left\{ 1, \frac{\pi(x^*)f(x_s | x)}{\pi(x_s)f(x^* | x)} \right\} \quad (4)$$

where x_s is the value of x on the parameter server prior to receiving the message.

As before, since the proposal distribution can depend on out-of-date states, Algorithm 6 cannot be directly analyzed in a Markov framework. Result 5 allows us to analyze convergence by pretending that communication is instantaneous, thereby allowing us to recover the Markov property and use it indirectly. Using this idea, proving that Algorithm 6 converges is straightforward: it essentially amounts to checking that the construction in Terenin et al. [11] can be modified to accommodate the needed communication pattern, which is close to immediate.

It is often possible to avoid evaluating the data-dependent terms $\pi(x^*)$, $\pi(x_s)$ on the parameter server. If we use, for example, a Hamiltonian Monte Carlo step, then $\pi(x)$ is immediately available from the worker node’s computation. This scheme directly addresses one of the disadvantages of Algorithm 2 by allowing each worker to have a copy of the last state it saw. It also allows the algorithm to complete the burn-in phase more quickly by rejecting unhelpful updates from out-of-date workers. Since our aim in this work is to showcase theory, we defer evaluation of performance to future work.

There are many other possible choices. Here, we have illustrated how analysis of Markov Chain Monte Carlo methods may be performed in the asynchronous distributed setting. We hope that future work discovers more such methods.

Acknowledgments

We are grateful to David Draper, Willie Neiswanger, and Hao Zhang for their thoughts.

References

- [1] G. M. Baudet. Asynchronous iterative methods for multiprocessors. *Journal of the Association for Computing Machinery*, 25(2):226–244, 1978. Cited on page 3.
- [2] D. P. Bertsekas. Distributed asynchronous computation of fixed points. *Mathematical Programming*, 27(1):107–120, 1983. Cited on page 3.
- [3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3(1):993–1022, 2003. Cited on page 1.
- [4] C. De Sa, K. Olukotun, and C. Ré. Ensuring rapid mixing and low bias for asynchronous Gibbs sampling. *arXiv:1602.07415*, 2016. Cited on page 2.
- [5] A. Frommer and D. B. Szyld. On asynchronous iterations. *Journal of Computational and Applied Mathematics*, 123(1):201–216, 2000. Cited on page 3.
- [6] Q. Ho, J. Cipar, H. Cui, S. Lee, J. K. Kim, P. B. Gibbons, G. A. Gibson, G. Ganger, and E. P. Xing. More effective distributed machine learning via a stale synchronous parallel parameter server. In *Advances in Neural Information Processing Systems*, pages 1223–1231, 2013. Cited on pages 1, 4.
- [7] A. Ihler and D. Newman. Understanding errors in approximate distributed Latent Dirichlet Allocation. *IEEE Transactions on Knowledge and Data Engineering*, 24(5):952–960, 2012. Cited on pages 1, 2.
- [8] M. Johnson, J. Saunderson, and A. Willsky. Analyzing Hogwild parallel Gaussian Gibbs sampling. In *Advances in Neural Information Processing Systems*, pages 2715–2723, 2013. Cited on page 2.
- [9] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, 1993. Cited on page 3.
- [10] D. Newman, A. Asuncion, P. Smyth, and M. Welling. Distributed algorithms for topic models. *Journal of Machine Learning Research*, 10:1801–1828, 2009. Cited on pages 1, 2.
- [11] A. Terenin, D. Simpson, and D. Draper. Asynchronous Gibbs Sampling. *arXiv:1509.08999*, 2016. Cited on pages 2–4, 6.
- [12] A. Terenin, S. Dong, and D. Draper. GPU-accelerated Gibbs Sampling: a case study of the Horseshoe Probit model. *arXiv:1608.04329*, 2016. Cited on page 1.
- [13] A. Terenin, M. Magnusson, L. Jonsson, and D. Draper. Pólya Urn Latent Dirichlet Allocation: a doubly sparse massively parallel sampler. *arXiv:1704.03581*, 2017. Cited on page 2.
- [14] J. Wei, W. Dai, A. Qiao, Q. Ho, H. Cui, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and E. P. Xing. Managed communication and consistency for fast data-parallel iterative analytics. In *Proceedings of the Sixth ACM Symposium on Cloud Computing*, pages 381–394, 2015. Cited on pages 1, 4.

Appendix A: proof of Theorem 4

Theorem 4. *Let P be a Markov operator with unique limiting distribution π . Let $x_0 \sim \mu_0$ for some arbitrary μ_0 . Let x_1, x_2, \dots be the sequence of values written to shared memory by the threads, and let μ_1, μ_2, \dots be their distributions. Assume that there are no partial reads or writes of x . Assume that P remains time-homogeneous if executed asynchronously. Assume that the maximum time between two updates made by the same thread is bounded above by b . Then we have that $\lim_{k \rightarrow \infty} \mu_k = \pi$.*

Proof. Fix an arbitrary μ_0 . Fix $b \in \mathbb{N}$. We assume the maximum time between two updates made by the same worker is bounded in the sense that we have $\mu_{k+1} = P(\mu_j)$ for some $k - b < j \leq k$: call this the *no worker dies* assumption. Define the values $d_k = \|\mu_k - \pi\|_{\text{TV}}$. Now, consider

$$d_k^* = \max\{d_j : k - b < j \leq k\} \quad (5)$$

defined for $j > b$. We claim that $d_k^* \geq d_{k+1}^*$. First, note that

$$d_{k+1}^* = \max\{d_j : k - b + 1 < j \leq k + 1\} = \max\{\max\{d_j : k - b + 1 < j \leq k\}, d_{k+1}\}. \quad (6)$$

Consider the left-hand term inside max. We have that

$$d_k^* = \max\{d_j : k - b < j \leq k\} \geq \max\{d_j : k - b + 1 < j \leq k\} \quad (7)$$

because max is taken over a smaller domain. Now, consider the right-hand term. By the *no worker dies* assumption and Result 3, we have that

$$d_k^* = \max\{d_j : k - b < j \leq k\} \geq d_{k+1}. \quad (8)$$

Together, this establishes that $d_k^* \geq d_{k+1}^*$. Since by definition $d_k^* \geq 0$, the sequence d_k^* is monotone and bounded, therefore it converges. We now claim that there exists a subsequence

$$d_{l_k}^* \text{ such that } d_{l_k}^* \rightarrow 0 \text{ as } k \rightarrow \infty. \quad (9)$$

For each value μ_k , let p_k denote the number of times the Markov transition kernel P was applied to obtain μ_k , i.e. if $\mu_k = P^3(\mu_0)$ then $p_k = 3$. Let

$$p_k^* = \min\{p_j : k - b < j \leq k\} \quad (10)$$

which is nondecreasing and unbounded by the *no worker dies* assumption – to see this formally, decompose min as in (6). Define the sequence of times

$$l_k = \arg \max\{p_j^* : 0 < j < k\} \quad (11)$$

which must also be unbounded. But then we have $P^{l_k+1}(\mu_0) = P[P^{l_k}(\mu_0)]$, so we can see by applying Definition 1 that as $k \rightarrow \infty$ we have

$$\|P^{l_k}(\mu_0) - \pi\| \rightarrow 0 \text{ and hence } d_{l_k}^* \rightarrow 0 \quad (12)$$

which gives the desired subsequence. Finally, since d_k^* is monotone, bounded below by 0, and admits an infinite subsequence that converges to 0, we have that $d_k^* \searrow 0$. Since $d_k^* \geq d_k$, this implies $d_k \rightarrow 0$, which completes the proof. ■

Appendix B: regularity conditions for Result 5

Result 5. Suppose that $H : E \rightarrow E$ is a weakly coupled Markov operator with unique coupled limiting distribution Π – these refer to the entire cluster, whereas π and \mathcal{M} refer to each worker. Let m be the number of workers, assumed fixed and finite. Assume the following.

$$\text{Box Condition: we have } E = \bigtimes_{i=1}^m \mathcal{M}.$$

Let H_i be the i th component of H with respect to E – this is each worker’s component of the coupled operator. Assume H is constructed from a Markov operator P such that it inherits Result 3 with respect to each H_i . Under appropriate regularity conditions given by Assumption 11 of Appendix B, H will converge asynchronously in the sense of Definition 9 of Appendix B. See also Remark 10.

Proof. Terenin et al. [11]. ■

Definition 9. Let $k \in \mathbb{N}_0$ be the total number of iterations performed by all workers. Let $s_i(k) \in \mathbb{N}_0$ be the total number of iterations on component i by all workers. Let I^k be an index set containing all components updated at iteration k . Let $\varepsilon^0 \in E$. Define the asynchronous iterations ε^k component-wise as follows.

$$\varepsilon_i^k = \begin{cases} H_i(\varepsilon_1^{s_1(k)}, \dots, \varepsilon_m^{s_m(k)}) & \text{if } i \in I^k, \\ \varepsilon_i^{k-1} & \text{otherwise.} \end{cases} \quad (13)$$

Note that this definition requires H to be time-homogeneous. We say that an asynchronous algorithm converges with respect to a given metric if ε^k converges to Π in that metric. For MCMC, we use the metric $\|\cdot\|_{TV}^\times = \sum_{i=1}^m \|\cdot\|_{TV}$ defined on E via the sum of the total variation metric on \mathcal{M} for each worker.

Remark 10. Result 5 also holds if we take that $E = \bigtimes_{i=1}^m \mathcal{M}_i$ for a sequence of measure spaces \mathcal{M}_i . For MCMC, we use $\mathcal{M}_i = \mathcal{M}$ to ensure $\varepsilon_i \rightarrow \pi$.

Assumption 11. Assume that all of the following hold.

No worker’s state is based on future values: $s_i(k) \leq k - 1$.

No worker stops permanently: $s_i(k) \rightarrow \infty$ as $k \rightarrow \infty$.

No component stops being updated or communicated: $|\{k \in \mathbb{N} : i \in I^k\}| = \infty$.