
Asynchronous Stochastic Gradient MCMC with Elastic Coupling

Jost Tobias Springenberg Aaron Klein Stefan Falkner Frank Hutter
Department of Computer Science, University of Freiburg
{springj,kleinaa,sfalkner,fh}@cs.uni-freiburg.de

Abstract

We consider parallel asynchronous Markov Chain Monte Carlo (MCMC) sampling for problems where we can leverage (stochastic) gradients to define continuous dynamics which explore the target distribution. We outline a solution strategy for this setting based on stochastic gradient Hamiltonian Monte Carlo sampling (SGHMC) which we alter to include an elastic coupling term that ties together multiple MCMC instances. The proposed strategy turns inherently sequential HMC algorithms into asynchronous parallel versions. First experiments empirically show that the resulting parallel sampler significantly speeds up exploration of the target distribution, when compared to standard SGHMC, and is less prone to the harmful effects of stale gradients than a naive parallelization approach.

1 Introduction and Background

Over the last years the ever increasing complexity of machine learning (ML) models together with the increasing amount of data that is available to train them has resulted in a great demand for parallel training and inference algorithms that can be deployed over many machines. To meet this demand, ML practitioners have increasingly relied on (asynchronous) parallel stochastic gradient descent (SGD) methods for optimizing model parameters [Recht et al., 2011, Dean et al., 2012, Zhang et al., 2015]. In contrast to this, the literature on efficient parallel methods for sampling from the posterior over model parameters given some data is much more scarce. Examples of algorithms for this setting typically are constrained to specific model classes [Ahmed et al., 2012, Ahn et al., 2015, Simsekli et al., 2015] or “only” consider data parallelism [Ahn et al., 2014]; in summary – with the exception of recent work by Chen et al. [2016] – a general sampling pendant of asynchronous SGD methods is missing.

In this paper we consider the general problem of sampling from an arbitrary posterior distribution over parameters $\theta \in \mathbb{R}^n$ given a set of observed datapoints $\mathbf{x} \in \mathcal{D}$ where we have K machines available to speed up the sampling process. Without loss of generality we can write the mentioned posterior as $p(\theta|\mathcal{D}) \propto \exp(-U(\theta))$ where we define $U(\theta)$ to be *the potential energy* $U(\theta) = -\sum_{\mathbf{x} \in \mathcal{D}} \log p(\mathbf{x}|\theta) - \log p(\theta)$. Many algorithms for solving this problem such as Hamiltonian Monte Carlo [Duane et al., 1987, Neal, 2010] further augment this potential energy with terms depending on auxiliary variables $\mathbf{y} \in \mathbb{R}^A$ to speed up sampling. In this more general case we write the posterior as $p(\mathbf{z}|\mathcal{D}) \propto \exp(-H(\mathbf{z}))$, where $\mathbf{z} = [\theta, \mathbf{y}] \in \mathbb{R}^m$ denotes the collection of all variables and $H(\mathbf{z})$ denotes the *Hamiltonian* $H(\mathbf{z}) = H(\theta, \mathbf{y}) = U(\theta) + g(\theta, \mathbf{y})$. Samples from $p(\theta|\mathcal{D})$ can then be obtained by sampling from $p(\mathbf{z}|\mathcal{D})$ and discarding \mathbf{y} if we additionally assume that marginalizing out \mathbf{y} from $p(\mathbf{z}|\mathcal{D})$ results only in a constant offset c ; that is we require $\int_{\mathbf{y} \in \mathbb{R}^A} \exp(-g(\theta, \mathbf{y})) d\mathbf{y} = c$.

1.1 Stochastic Gradient Hamiltonian Monte Carlo

Stochastic gradient MCMC (SGMCMC) algorithms solve the above described sampling problem by assuming access to a – possibly noisy – gradient of the potential $U(\theta)$ with respect to parameters θ . In this case – assuming properly chosen auxiliary variables – sampling can be performed via an analogy to physics by simulating a system based on the Hamiltonian $H(\mathbf{z})$. More precisely, following the general formulation from Ma et al. [2015], one can simulate a stochastic differential equation (SDE) of the form

$$d\mathbf{z} = f(\mathbf{z}) dt + \sqrt{2D(\mathbf{z})} d\mathbf{w}_t, \quad (1)$$

where $f(\mathbf{z}) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ denotes the deterministic drift incurred by $H(\mathbf{z})$, $D(\mathbf{z}) : \mathbb{R}^m \rightarrow \mathbb{R}^m \times \mathbb{R}^m$ is a diffusion matrix, the square root is applied element-wise, and $\mathbf{w}_t \in \mathbb{R}^m$ denotes Brownian motion. Under fairly general assumptions on the functions f and D the unique stationary distribution of this system is equivalent to the posterior $p(\theta|\mathcal{D})$. Specifically, Ma et al. [2015] showed that if $f(\mathbf{z})$ is of the following specialized form (in which we are free to choose $D(\mathbf{z})$ and $Q(\mathbf{z})$):

$$f(\mathbf{z}) = -(D(\mathbf{z})+Q(\mathbf{z})) \begin{pmatrix} \nabla_{\theta}U(\theta) + \nabla_{\theta}g(\theta, \mathbf{y}) \\ \nabla_{\mathbf{y}}g(\theta, \mathbf{y}) \end{pmatrix} + \Gamma(\mathbf{z}), \quad \Gamma_i(\mathbf{z}) = \sum_{j=1}^m \frac{\partial}{\partial \mathbf{z}_j} (D_{i,j}(\mathbf{z}) + Q_{i,j}(\mathbf{z})), \quad (2)$$

then the stationary distribution is equivalent to the posterior if $D(\mathbf{z})$ is positive semi-definite and $Q(\mathbf{z})$ is skew-symmetric. Importantly, this holds also if only noisy estimates $\nabla_{\theta}\tilde{U}(\theta)$ of the gradient $\nabla_{\theta}U(\theta)$ computed on a randomly sampled subset of the data are available (as in our case).

1.1.1 Practical SGMCMC implementations

In practice, for any choice of $H(\mathbf{z})$, $D(\mathbf{z})$ and $Q(\mathbf{z})$, simulating the differential equation (1) on a digital computer involves two approximations. First, the SDE is simulated in discretized steps resulting in the update rule

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \epsilon_t \left[(D(\mathbf{z}) + Q(\mathbf{z})) \begin{pmatrix} \nabla_{\theta}U(\theta) - \nabla_{\theta}g(\theta, \mathbf{y}) \\ \nabla_{\mathbf{y}}g(\theta, \mathbf{y}) \end{pmatrix} - \Gamma(\mathbf{z}) \right] + \mathcal{N}(0, 2\epsilon_t D(\mathbf{z}_t)), \quad (3)$$

where we slightly abuse notation and take $\mathcal{N}(\mu, \Sigma)$ to denote the addition of a *sample from* an m-dimensional multivariate Gaussian distribution. Second, when dealing with large datasets, exact computation of the gradient $\nabla_{\theta}U(\theta)$ becomes computationally prohibitive and one thus relies on a stochastic approximation computed on a randomly sampled subset of the data: $\nabla_{\theta}\tilde{U}(\theta)$ with $\tilde{U}(\mathbf{z}) = \frac{N}{|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} \log p(\mathbf{x}|\theta) - \log p(\theta)$ where $\mathcal{B} \subset \mathcal{D}$. The stochastic gradient $\nabla_{\theta}\tilde{U}(\theta)$ is then Gaussian distributed with some variance V ; leading to the noise term in the above described SDE. Using these two approximations one can derive the following discretized system of equations for a stochastic gradient variant of Hamiltonian Monte Carlo

$$\begin{aligned} \theta_{t+1} &= \theta_t + \epsilon \mathbf{M}^{-1} \mathbf{p}_t \\ \mathbf{p}_{t+1} &= \mathbf{p}_t - \epsilon \nabla_{\theta}\tilde{U}(\theta_t) - \epsilon \mathbf{V} \mathbf{M}^{-1} \mathbf{p}_t + \mathcal{N}(0, 2\epsilon \mathbf{V}), \end{aligned} \quad (4)$$

which, following Ma et al. [2015], can be seen as an instance of Equation (3) with $\mathbf{y} = \mathbf{p}$, $g(\theta, \mathbf{p}) = \mathbf{p}^T \mathbf{M}^{-1} \mathbf{p}$ and where $D(\mathbf{z}) = \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{V} \end{pmatrix}$ and $Q(\mathbf{z}) = \begin{pmatrix} 0 & \mathbf{I} \\ -\mathbf{I} & \mathbf{V} \end{pmatrix}$.

2 Parallelization schemes for SG-MCMC

We now show how one can utilize the computational power of K machines to speed up a given sampling procedure relying on the dynamics described in Equations (4). As mentioned before, the update equations derived in Section 1.1.1 involve alternating updates to both the parameters and the auxiliary variables, leading to an inherently sequential algorithm. If we now want to utilize K machines to speed up sampling we thus face the non-trivial challenge of parallelizing these updates. In general we can identify two solutions¹:

¹We note that if the computation of the stochastic gradient $\nabla_{\theta}\tilde{U}(\theta)$ is based on a large number of data-points (or if we want to reduce the variance of our estimate by increasing $|\mathcal{B}|$) we could potentially spread out the

I) For a naive parallelization strategy we can send the variables θ to K different machines every s steps from a parameter server. Each machine then computes a gradient estimate for the current step $\nabla_{\theta} \tilde{U}(\tilde{\theta}_t^k)$ (note that we only approximately have $\tilde{\theta}_t^k \approx \theta_t$ due to the communication period s , i.e. $\tilde{\theta}_t^k$ at each machine might be a stale parameter). The server then waits for O gradient estimates $\nabla_{\theta} \tilde{U}(\tilde{\theta}_t^k)$ to be sent back and simulates the system from Eq. (1) using $\nabla_{\theta} \tilde{U}(\theta_t) \approx \frac{1}{O} \sum_{k=1}^O \nabla_{\theta} \tilde{U}(\tilde{\theta}_t^k)$;

II) We can set-up K MCMC chains (one per machine) which independently update a parameter vector \mathbf{z}^k (where $k \in [1, K]$ denotes the machine), following the dynamics from Eq. (4).

While the second approach clearly results in Markov chains that asymptotically sample from the correct distribution – and might result in a more diverse set of samples than a simulation using a single chain – it is also clear that it cannot speed up convergence of the individual chains (our desired goal) as there is no interaction between them. The first approach, on the other hand, is harder to analyze. We can observe that if $s = 1$ and we wait for $O = K$ gradient estimates in each step we obtain an SG-MCMC algorithm with parallel gradient estimates but synchronous updates of the dynamic equations. Consequently, such a setup preserves the guarantees of standard SG-MCMC but *requires synchronization between all machines in each step*. In a real-world experiment (where we might have heterogeneous machines and communication delays) this will result in a large communication overhead. For choices of $s > 1$ and $O < K$ – the regime we are interested in – we cannot rely on the standard convergence analysis for SG-MCMC anymore. Nonetheless, if we concentrate on the analysis of $s > 1$ and $O = 1$ (i.e. completely asynchronous updates) we can interpret the stale parameters to simply result in more noisy estimates of $\nabla_{\theta} \tilde{U}(\theta_t)$ that can be used within the dynamic equations from Eq. 4. The efficacy of such a parallelization scheme then intuitively depends on the amount of additional noise introduced by the stale parameters and requires a new convergence analysis. During the preparation of this manuscript, concurrent work on SGMCMC with stale gradients derived a theoretical foundation for this intuition [Chen et al., 2016]. Interestingly, we will in the following empirically show that the additional noise is unproblematic for small s in the range $1 < s < 4$ (for which a convergence speed-up with naive parallelization can therefore be achieved, but which result in a large communication overhead in distributed systems) but becomes problematic with growing s . We believe these results are in accordance with the mentioned recent work by Chen et al. [2016], yet a unification of their theory with our proposed new algorithm remains as important future work.

3 Stochastic Gradient MCMC with Elastic Coupling

Given the negative analysis from Section 2 one might wonder whether approach **II)** (the idea of running K parallel MCMC chains) can be altered such that the K chains are only loosely coupled; allowing for faster convergence while avoiding excessive communication. To this end we propose to consider the following alternative parallelization scheme:

IIa) To speed up K SGMCMC chains we couple the K parameter vectors through an additional center variable \mathbf{c} to which they are elastically attached. We collect updates to this center variable at a central server and broadcast an updated version of it every s steps across all machines. We note that an approach based on this idea was recently utilized to derive an asynchronous SGD optimizer in Zhang et al. [2015], serving as our main inspiration. A discussion of the connection between their deterministic and our stochastic dynamics is presented in Section 5.

To derive an asynchronous SGMCMC variant with K samplers and elastic coupling – as described above – we consider an augmented Hamiltonian with $\mathbf{z} = [\theta^1, \dots, \theta^K, \mathbf{p}^1, \dots, \mathbf{p}^K, \mathbf{c}, \mathbf{r}]$:

$$H(\mathbf{z}) = \sum_{i=1}^K \left(U(\theta^i) + \mathbf{p}^{iT} \mathbf{M}^{-1} \mathbf{p}^i \right) + \frac{1}{K} \sum_{i=1}^K \frac{\alpha}{2} \|\theta^i - \mathbf{c}\|_2^2 + \mathbf{r}^T \mathbf{M}^{-1} \mathbf{r}, \quad (5)$$

where we can interpret \mathbf{c} as a centering mass (with momentum \mathbf{r}) through which the asynchronous samplers are elastically coupled and α specifies the coupling strength. It is easy to see that for

data-set \mathcal{D} over the K machines allowing us to parallelize computation without the need for asynchronous updating. While this is an interesting problem in its own right and there already exists a considerable amount of literature on running parallel MCMC over sub-sets of data Scott et al. [2016], Rabinovich et al. [2015], Neiswanger et al. [2014] we here focus on parallelization schemes that do not make this assumption because of their broad applicability.

$\alpha = 0$ we can decompose the sum from Eq. (5) into K independent terms which each constitute a Hamiltonian corresponding to the standard SGHMC case (and we thus recover the setup of K independent MCMC chains). Further, for $\alpha > 0$ we obtain a joint, altered, Hamiltonian in which – as desired – the K parameter vectors are elastically coupled. Simulating a system according to this Hamiltonian exactly would again result in an algorithm requiring synchronization in each simulation step (since the change in momentum for all parameters depends on the position of the center variable and thus, implicitly, on all other parameters).

If we, however, assume only a noisy measurement of the center variable and its momentum is available in each step we can derive a mostly asynchronous algorithm for updating each θ^i . To achieve this let us assume we store our current estimate for the center variable and its momentum at a central server. This server receives updates for \mathbf{c} and \mathbf{r} from each of the K samplers every s steps and replies with their current values. Assuming a Gaussian normal distribution on the error of this current value each sampler then keeps track of a noisy estimate $\tilde{\mathbf{c}} \approx \mathbf{c}$ of the center variable which is used for simulating updates derived from the Hamiltonian in Equation (5). From these assumptions we derive the following discretized dynamical equations:

$$\begin{aligned}\theta_{t+1}^i &= \theta_t^i + \epsilon \mathbf{M}^{-1} \mathbf{p}_t^i, \\ \mathbf{c}_{t+1} &= \mathbf{c}_t + \epsilon \mathbf{M}^{-1} \mathbf{r}_t, \\ \mathbf{p}_{t+1}^i &= \mathbf{p}_t^i - \epsilon \nabla \tilde{U}(\theta_t^i) - \epsilon \mathbf{V} \mathbf{M}^{-1} \mathbf{p}_t^i - \epsilon \alpha (\theta_t^i - \mathbf{c}_t) + \mathcal{N}(0, 2\epsilon^2(\mathbf{V} + \mathbf{C})), \\ \mathbf{r}_{t+1} &= \mathbf{r}_t - \epsilon \mathbf{C} \mathbf{M}^{-1} \mathbf{r}_t - \epsilon \alpha \frac{1}{K} \sum_{i=1}^K (\mathbf{c}_t - \theta_t^i) + \mathcal{N}(0, 2\epsilon^2 \mathbf{C}),\end{aligned}\tag{6}$$

where \mathbf{V} specifies the noise due to stochastic gradient estimates and \mathbf{C} is the variance of the aforementioned noisy center variable. As before, we use the notation $\mathcal{N}(\mu, \Sigma)$ to refer to a sample from a Gaussian distribution with mean μ and covariance Σ . We note that, while the presented dynamical equations were derived from SGHMC, the elastic coupling idea does not depend on the basic Hamiltonian from Equation (4). We can thus derive similar asynchronous samplers for any SGMCMC variant including first order stochastic Langevin dynamics [Welling and Teh, 2011] or any of the more advanced techniques reviewed in Ma et al. [2015].

When inspecting the Equations (6) we can observe that, similar to approach **I**, they also contain an additional noise source: noise is injected into the system due to potential staleness of the center variable. However, since this noise only indirectly affects the parameters θ^i one might hope that the center variable acts as a buffer, damping the injected noise. If this were the case, we would expect the resulting algorithm to be more robust to communication delays than the naive parallelization approach. In addition to this consideration, the proposed dynamical equations have a convenient form which makes it easy to verify that they fulfill the conditions for a valid SGMCMC procedure.

Proposition 3.1 *The dynamics of the system from Eq. (6) has the posterior distribution $p(\theta \mid \mathcal{D})$ as the stationary distribution for all K samplers.*

Proof To show this, we first establish that Equations (6) correspond to a discretized dynamics following the general form given by Equations (1) and (2) where $D(\mathbf{z}) = \text{diag}([0, \mathbf{V}, 0, \mathbf{C}])$ and $Q(\mathbf{z}) = \begin{pmatrix} A & B \\ B & A \end{pmatrix}$, with $A = \begin{pmatrix} 0 & \mathbf{I} \\ -\mathbf{I} & 0 \end{pmatrix}$ and $B = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$, thus fulfilling the requirements outlined in Section 1.1. Furthermore, to see that marginalization of the auxiliary variables results in a constant offset, we can first identify $g(\theta^i, \mathbf{y}) = \alpha/2K \|\theta^i - \mathbf{c}\|_2^2 + \mathbf{r}^T \mathbf{M}^{-1} \mathbf{r} + \mathbf{p}^{iT} \mathbf{M}^{-1} \mathbf{p}^i$, with $\mathbf{y} = [\mathbf{c}, \mathbf{r}, \mathbf{p}]$. Solving the integral $\int_{\mathbf{y} \in \mathbb{R}^A} \exp(-g(\theta^i, \mathbf{y})) d\mathbf{y}$ then amounts to evaluating Gaussian integrals and is therefore easily checked to be constant, as required. Thus, simulating the the dynamical equations results in samples from $p(\theta, \mathbf{y} \mid \mathcal{D})$ and discarding the auxilliary variables \mathbf{y} gives the desired result.

4 First Experiments

We designed an initial set of three experiments to validate our sampler. First, to get an intuition for the behavior of the elastic coupling term we tested the sampler on a low dimensional toy example. In Figure 1 we show the first 100 steps taken by both standard SGHMC (left) and our elastically coupled

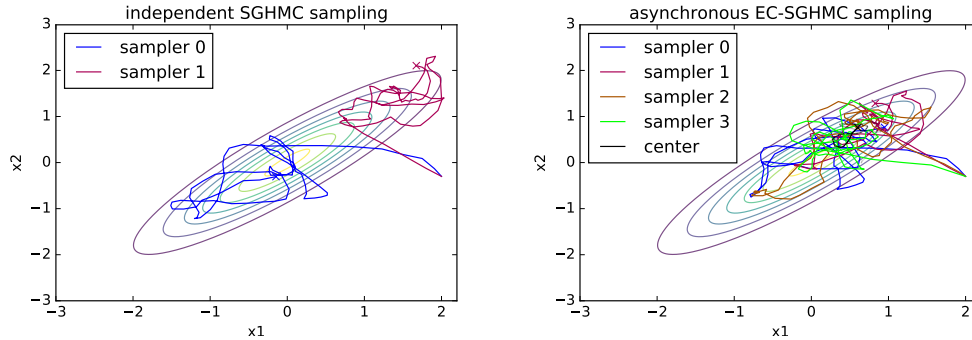


Figure 1: Comparison between the first 100 sampling steps performed by SGHMC and the elastically coupled SGHMC variant for sampling from a simple two dimensional Gaussian distribution. An animated video of the samplers can be found at <https://goo.gl/ZZv1fG>.

sampler (EC-SGHMC) with four parallel threads (right) when sampling from a two-dimensional Gaussian distribution (starting from the same initial guess). The hyperparameters were set to $\alpha = 1$, $\epsilon = 1e - 2$, $\mathbf{C} = \mathbf{V} = \mathbf{I}$. We observe that two independent runs of SGHMC take fairly different initial paths and, depending on the noise, it can happen that SGHMC only explores low-density regions of the distribution in its first steps (cf. purple curve). In contrast the four samplers with elastic coupling quickly sample from high density regions and show coherent behaviour.

As a second experiment, we compare EC-SGHMC to standard SGHMC and the naive parallelization described in Section 2 (Async SGHMC in the following). We use each method for sampling the weights of a two layer fully connected neural network (800 units per layer, ReLU activations, Gaussian prior on the weights, batch size 100) that is applied to classifying the MNIST dataset. For our purposes, we can interpret the neural network as a function that parameterizes a probability distribution over classes

$$p(y = i | \mathbf{x}, \theta) \propto \exp(f_i(\mathbf{x}, \theta)), \quad (7)$$

where y is the label of the given image and $f(\mathbf{x}, \theta)$ is the output vector of the neural network with parameters θ . We place a Gaussian prior on the network weights $p(\theta) \propto \exp(-\lambda \|\theta\|_2^2)$ (where we chose $\lambda = 10^{-5}$) and sample from the posterior

$$p(\theta | \mathcal{D}) \propto \prod_{(\mathbf{x}^j, y^j) \in \mathcal{D}} p(y^j | \mathbf{x}^j, \theta) p(\theta), \quad (8)$$

for a given dataset \mathcal{D} . The results of this experiment are depicted in Figure 2 (left). We plot the negative log likelihood over time and observe that both parallel samplers (using $K = 6$ parallel threads) perform significantly better than standard SGHMC. However, when we increase the communication delay and only synchronize threads every $s = 8$ steps the additional noise injected into the sampler via this procedure becomes problematic for Async SGHMC whereas EC-SGHMC copes with it much more gracefully.

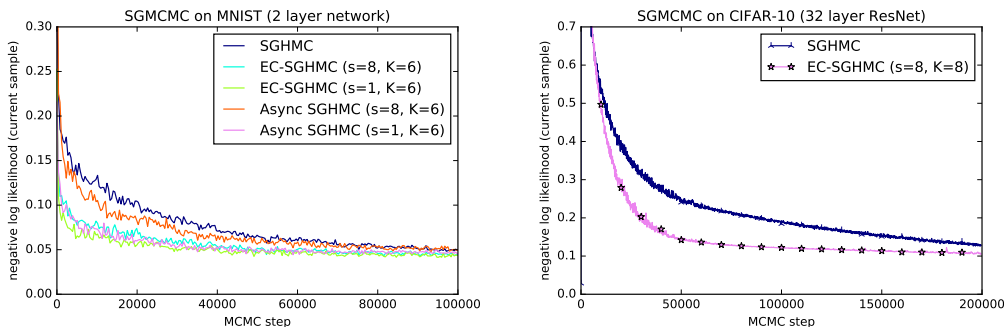


Figure 2: Comparison between different SGMCMC samplers for sampling from the posterior over neural network weights for a fully connected network on MNIST (left) and a residual network on CIFAR-10 (right). Best viewed in color.

Finally, to test the scalability of our approach, we sampled the weights of a 32-layer residual net applied to the CIFAR-10 dataset. We again aim to sample the posterior given by Equation (8) only now the neural network $f(\mathbf{x}, \theta)$ is the 32-layer residual network described in He et al. [2016] (with batch-normalization removed). The results of this experiment are depicted in Figure 2 (right) showing that, again, EC-SGHMC leads to a significant speed-up over standard SGHMC sampling.

5 Connection to Elastic Averaging SGD

As described in Section 3 an elastic coupling technique similar to the one used in this manuscript was recently used to accelerate asynchronous stochastic gradient descent [Zhang et al., 2015]. Although the purpose of our paper is not to derive a new SGD method – we instead aim to arrive at a scalable MCMC sampler – it is instructive to take a closer look at the connection between the two methods.

To establish this connection we aim to re-derive the elastic averaging SGD (EASGD) method from Zhang et al. [2015] as the deterministic limit of the dynamical equations from (6). Removing the added noise from these equations, setting \mathbf{M} to the identity matrix, and performing the variable substitutions $\mathbf{v}^i = \epsilon \mathbf{M} \mathbf{p}^i$, $\mathbf{h}^i = \epsilon \mathbf{M} \mathbf{r}^i$, $\xi = \mathbf{V} = \mathbf{C}$ yields the dynamical equations

$$\begin{aligned}\theta_{t+1}^i &= \theta_t^i + \mathbf{v}_t^i, \\ \mathbf{c}_{t+1} &= \mathbf{c}_t + \mathbf{h}_t, \\ \mathbf{v}_{t+1}^i &= \mathbf{p}_t^i - \epsilon \nabla \tilde{U}(\theta_t^i) - \xi \mathbf{v}_t^i - \epsilon \alpha (\theta_t^i - \mathbf{c}_t), \\ \mathbf{h}_{t+1} &= \mathbf{r}_t - \xi \mathbf{h}_t - \epsilon \alpha \frac{1}{K} \sum_{i=1}^K (\mathbf{c}_t - \theta_t^i).\end{aligned}\tag{9}$$

In comparison, re-writing the updates for the EASGD variant with momentum (EAMSGD) from Zhang et al. [2015] in our notation (and replacing the Nesterov momentum with a standard momentum) we obtain

$$\begin{aligned}\theta_{t+1}^i &= \theta_t^i + \mathbf{v}_t^i - \epsilon \alpha (\theta_t^i - \mathbf{c}_t), \\ \mathbf{c}_{t+1} &= \mathbf{c}_t - \epsilon \alpha \frac{1}{K} \sum_{i=1}^K (\mathbf{c}_t - \theta_t^i), \\ \mathbf{v}_{t+1}^i &= \mathbf{p}_t^i - \epsilon \nabla \tilde{U}(\theta_t^i) - \xi \mathbf{v}_t^i,\end{aligned}\tag{10}$$

where, additionally, Zhang et al. [2015] propose to only update \mathbf{c}_{t+1} every s steps and drop the terms including \mathbf{c}_{t+1} in all update equations in intermittent steps.

As expected, at first glance the two sets of update equations look very similar. Interestingly, they do however differ with respect to integration of the elastic coupling term and the updates to the center variable: In the EAMSGD Equations (10) the center variables are not augmented with a momentum term and the elastic coupling force influences the parameter values θ directly rather than, indirectly, through their momentum \mathbf{p} .

From the physics perspective that we adopt in this paper these updates are thus “wrong” in the sense that they break the interpretation of the variables θ , \mathbf{c} and \mathbf{p} as generalized coordinates and generalized momenta. It should be noted that there also is no straight-forward way to recover a valid SGMCMC sampler corresponding to a stochastic variant of Equations (10) from the Hamiltonian given in Equation (5). Our derivation thus suggests alternative update equations for EAMSGD. An interesting avenue for future experiments thus is to thoroughly compare the deterministic updates from Equations (9) with the EAMSGD updates both in terms of empirical performance and with respect to their convergence properties. An initial test we performed suggests that the former perform at least as good as EAMSGD. We also note that EASGD without momentum can exactly be recovered as the deterministic limit of our approach (without the above described discrepancies) if we were to randomly re-sample the auxiliary momentum variables in each step – and would thus simulate stochastic gradient Langevin dynamics Welling and Teh [2011].

6 Conclusion

In this manuscript we have considered the problem of parallel asynchronous MCMC sampling with stochastic gradients. We have introduced a new algorithm for this problem based on the idea of

elastically coupling multiple SGMCMC chains. First experiments suggest that the proposed method compares favorably to a naive parallelization strategy but additional experiments are required to paint a conclusive picture. We have further discussed the connection between our method and the recently proposed stochastic averaging SGD (EASGD) optimizer from Zhang et al. [2015], revealing an alternative variant of EASGD with momentum.

References

- Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Proc. of NIPS'11*, 2011.
- Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc'auelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc V. Le, and Andrew Y. Ng. Large scale distributed deep networks. In *Proc. of NIPS'12*, 2012.
- Sixin Zhang, Anna E. Choromanska, and Yann LeCun. Deep learning with elastic averaging SGD. In *Proc. of NIPS'15*, 2015.
- Amr Ahmed, Moahmed Aly, Joseph Gonzalez, Shравan Narayanamurthy, and Alexander J. Smola. Scalable inference in latent variable models. In *Proc. of WSDM '12*, 2012.
- Sungjin Ahn, Anoop Korattikara, Nathan Liu, Suju Rajan, and Max Welling. Large-scale distributed Bayesian matrix factorization using stochastic gradient MCMC. In *Proc. of KDD'15*, 2015.
- Umut Simsekli, Hazal Koptagel, Hakan Güldas, A. Taylan Cemgil, Figen Öztoprak, and S. Ilker Birbil. Parallel stochastic gradient Markov Chain Monte Carlo for matrix factorisation models. In *arxiv:1506.01418*, 2105.
- Sungjin Ahn, Babak Shahbaba, and Max Welling. Distributed stochastic gradient MCMC. In *Proc. of ICML'14*, 2014.
- Changyou Chen, Nan Ding, Chunyuan Li, Yizhe Zhang, and Lawrence Carin. Stochastic gradient MCMC with stale gradients. In *Proc. of NIPS'16*, 2016.
- Simon Duane, Anthony D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 1987.
- Radford M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, pages 113–162, 2010.
- Yi-An Ma, Tuanqi Chen, and Emily B. Fox. A complete recipe for stochastic gradient MCMC. In *Proc. of NIPS'15*, 2015.
- Steven L. Scott, Alexander W. Blocker, and Fernando V. Bonassi. Bayes and big data: The consensus monte carlo algorithm. *International Journal of Management Science and Engineering Management*, 2016.
- Maxim Rabinovich, Elaine Angelino, and Michael I Jordan. Variational consensus monte carlo. In *Proc. of NIPS*, 2015.
- Willie Neiswanger, Chong Wang, and Eric P. Xing. Asymptotically exact, embarrassingly parallel MCMC. In *Proc. of UAI*, 2014.
- Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proc. of ICML*, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of CVPR'16*, 2016.