
Exponential Stochastic Cellular Automata for Massively Parallel Inference

Manzil Zaheer*
Carnegie Mellon

Michael Wick
Oracle Labs

Jean-Baptiste Tristan†
Oracle Labs

Alex Smola
Carnegie Mellon

Guy L Steele Jr
Oracle Labs

Abstract

We propose an embarrassingly parallel, memory efficient inference algorithm for latent variable models in which the complete data likelihood is in the exponential family. The algorithm is a stochastic cellular automaton and converges to a valid *maximum a posteriori* fixed point. Applied to latent Dirichlet allocation we find that our algorithm is over an order or magnitude faster than the fastest current approaches. A simple C++/MPI implementation on a 20-node Amazon EC2 cluster samples at more than **1 billion tokens per second**. We process 3 billion documents and achieve predictive power competitive with collapsed Gibbs sampling and variational inference.

1 Introduction

In the past decade, frameworks such as stochastic gradient descent (SGD) [28] and map-reduce [8] have enabled machine learning algorithms to scale to larger and large datasets. However, these frameworks are not always applicable to Bayesian latent variable models with rich statistical dependencies and intractable gradients. Variational methods [14] and Markov chain Monte-Carlo (MCMC) [10] have thus become the *sine qua non* for inferring the posterior in these models.

Sometimes—due to the concentration of measure phenomenon associated with large sample sizes—computing the full posterior is unnecessary and *maximum a posteriori* (MAP) estimates suffice. It is hence tempting to employ gradient descent. However, for latent variable models such as latent Dirichlet allocation (LDA), calculating gradients involves expensive expectations over rich sets of variables [26].

MCMC is an appealing alternative, but traditional algorithms such as the Gibbs sampler are inherently sequential and the extent to which they can be parallelized depends heavily upon how the structure of the statistical model interacts with the data. For instance, chromatic sampling [11] is infeasible for LDA, due to its dependence structure. We propose an alternate approach based on stochastic cellular automata (SCA). The automaton is massively parallel like conventional cellular automata, but employs stochastic updates.

Our proposed algorithm, exponential SCA (ESCA), is a specific way of mapping inference in latent variable models with complete data likelihood in the exponential family into an instance of SCA. ESCA has a minimal memory footprint because it stores only the data and the minimal sufficient statistics (by the very definition of sufficient statistics, the footprint cannot be further reduced). In contrast, variational approaches such as stochastic variational inference (SVI) [13] require storing the variational parameters, while MCMC-based methods, such as YahooLDA [30] require storing the latent variables. Thus, ESCA substantially reduces memory costs, enabling larger datasets to fit in memory, and significantly reducing communication costs in distributed environments.

Furthermore, the sufficient statistics dramatically improves efficiency. Typically, in the general case of SCA, updating a cell requires first assembling the values of all the neighboring cells before aggregating them into a local stochastic update. However, in ESCA, the sufficient statistics adequately summarize the states of the neighbors; the computational load is therefore small and perfectly balanced across the cells.

We demonstrate how ESCA is a flexible framework for exponential latent variable models such as LDA. In our experiments, we process over 3 billion documents at a rate of 570 million tokens per second on a small cluster of 4 commodity servers. That said, ESCA is much more general. Table 1 explicitly lists some of the more common modeling choices for which ESCA can

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

*Work done when an intern at Oracle Labs.

†Corresponding author.

Table 1: Examples of some popular models to which ESCA is applicable.

mix. component/emitter	Bernoulli	Multinomial	Gaussian	Poisson
Categorical	Latent Class Model [9]	Unigram Document Clustering	Mixture of Gaussians [21]	
Dirichlet Mixture	Grade of Membership Model [35]	Latent Dirichlet Allocation [2]	Gaussian-LDA [6]	GaP Model [3]

be easily employed. Our algorithm implicitly simulates stochastic expectation maximization (SEM), and is thus provably correct in the sense that it converges in distribution to a stationary point of the *posterior*.

2 Exponential SCA

Stochastic cellular automata (SCA), also known as probabilistic cellular automata, or locally-interacting Markov chains, are a stochastic version of a discrete-time, discrete-space dynamical system in which a noisy local update rule is homogeneously and synchronously applied to every site of a discrete space. They have been studied in statistical physics, mathematics, and computer science, and some progress has been made toward understanding their ergodicity and equilibrium properties. A recent survey [18] is an excellent introduction to the subject, and a dissertation [17] contains a comprehensive and precise presentation of SCA.

The automaton, as a (stochastic) discrete-time, discrete-space dynamical system, is given by an evolution function $\Phi : \mathcal{S} \rightarrow \mathcal{S}$ over the state space $\mathcal{S} = \mathcal{Z} \rightarrow \mathcal{C}$ which is a mapping from the space of cell identifiers Z to cell values C . The global evolution function applies a local function $\phi_{\mathfrak{z}}(c_1, c_2, \dots, c_r) \mapsto c$ s.t. $c_i = s(\mathfrak{z}_i)$ to every cell $\mathfrak{z} \in \mathcal{Z}$. That is, ϕ examines the values of each of the neighbors of cell \mathfrak{z} and then stochastically computes a new value c . The dynamics begin with a state $s_0 \in \mathcal{S}$ that can be configured using the data or a heuristic.

Exponential SCA (ESCA) is based on SCA but achieves better computational efficiency by exploiting the structure of the sufficient statistics for latent variable models in which the complete data likelihood is in the exponential family. Most importantly, the local update function ϕ for each cell depends only upon the sufficient statistics and thus does *not* scale linearly with the number of neighbors.

2.1 Latent Variable Exponential Family

Latent variable models are useful when reasoning about partially observed data such as collections of text or images in which each *i.i.d.* data point is a document or image. Since the same local model is applied to each data point, they have the following form

$$p(\mathbf{z}, \mathbf{x}, \eta) = p(\eta) \prod_i p(z_i, x_i | \eta). \tag{1}$$

Our goal is to obtain a MAP estimate for the parameters η that explain the data \mathbf{x} through the latent variables \mathbf{z} . To expose maximum parallelism, we want each cell in the automaton to correspond to a data point and its latent variable. However, in general all latent variables depend on each other via the global parameters η , and thus the local evolution function ϕ would have to examine the values of every cell in the automaton.

Fortunately, if we further suppose that the complete data likelihood is in the exponential family, i.e.,

$$p(z_i, x_i | \eta) = \exp(\langle T(z_i, x_i), \eta \rangle - g(\eta)) \tag{2}$$

then the complete and sufficient statistics are given by

$$T(\mathbf{z}, \mathbf{x}) = \sum_i T(z_i, x_i) \tag{3}$$

and we can thus express any estimator of interest as a function of just $T(\mathbf{z}, \mathbf{x})$. Further, when employing expectation maximization (EM), the M-step is possible in closed form for many members of the exponential family. This allows us to reformulate the cell level updates to depend only upon the sufficient statistics instead of the neighboring cells. The idea is that, unlike SCA (or MCMC in general) which produces a sequence of states that correspond to complete variable assignments s^0, s^1, \dots via a transition kernel $q(s^{t+1} | s^t)$, ESCA produces a sequence of sufficient statistics T^0, T^1, \dots directly via an evolution function $\Phi(T^t) \mapsto T^{t+1}$.

2.2 Stochastic EM

Before we present ESCA, we must first describe stochastic EM (SEM). Suppose we want the MAP estimate for η and employ a traditional expectation maximization (EM) approach:

$$\max_{\eta} p(\mathbf{x}, \eta) = \max_{\eta} \int p(\mathbf{z}, \mathbf{x}, \eta) \mu(d\mathbf{z})$$

EM finds a mode of $p(\mathbf{x}, \eta)$ by iterating two steps:

E-step Compute in parallel $p(z_i | x_i, \eta^{(t)})$.

M-step Find $\eta^{(t+1)}$ that maximizes the expected value of the log-likelihood with respect to the conditional probability, i.e.

$$\begin{aligned} \eta^{(t+1)} &= \arg \max_{\eta} \mathbb{E}_{\mathbf{z} | \mathbf{x}, \eta^{(t)}} [\log p(\mathbf{z}, \mathbf{x}, \eta)] \\ &= \xi^{-1} \left(\frac{1}{n + n_0} \sum_i \mathbb{E}_{\mathbf{z} | \mathbf{x}, \eta^{(t)}} [T(z_i, x_i)] + T_0 \right) \end{aligned}$$

where $\xi(\eta) = \nabla g(\eta)$ is invertible as $\nabla^2 g(\theta) \succ 0$ and n_0, T_0 parametrize the conjugate prior.¹

Although EM exposes substantial parallelism, it is difficult to scale, since the dense structure $p(z_i|x_i, \eta^{(t)})$ defines values for all possible outcomes for z and thus puts tremendous pressure on memory bandwidth.

To overcome this we introduce sparsity by employing stochastic EM (SEM) [4]. SEM substitutes the E-step for an S-step that replaces the full distribution with a single sample:

S-step Sample $z_i^{(t)} \sim p(z_i|x_i; \eta^{(t)})$ in parallel.

Subsequently, we perform the M-step using the imputed data instead of the expectation. This simple modification overcomes the computational drawbacks of EM for cases in which sampling from $p(z_i|x_i; \eta^{(t)})$ is feasible. We can now employ fast samplers, such as the alias method, exploit sparsity, reduce CPU-RAM bandwidth while still maintaining massive parallelism.

The S-step has other important consequences. Notice that the M-step is now a simple function of current sufficient statistics. This implies that the conditional distribution for the next S-step is expressible in terms of the complete sufficient statistics

$$p(z_i|x_i; \eta^{(t)}) = f(z_i, T(\mathbf{x}, \mathbf{z}^{(t)})).$$

Thus each S-step depends only upon the sufficient statistics generated by the previous step. Therefore, we can operate directly on sufficient statistics without the need to assign or store latent variables/states. Moreover it opens up avenues for distributed and parallel implementations that execute on an SCA.

2.3 ESCA for Latent Variable Models

SEM produces an alternating sequence of S and M steps in which the M step produces the parameters necessary for the next S step. Since we can compute these parameters on the fly there is no need for an explicit M step. Instead, ESCA produces a sequence consisting only of S steps. We require the exponential family to ensure that these steps are both efficient and compact. We now present ESCA more formally.

Define an SCA over the state space \mathcal{S} of the form:

$$\mathcal{S} = \mathcal{Z} \rightarrow \mathcal{K} \times \mathcal{X} \tag{4}$$

where \mathcal{Z} is the set of cell identifiers (e.g., one per token in a text corpus), \mathcal{K} is the domain of latent variables, and \mathcal{X} is the domain of the observed data.

The initial state s_0 is the map defined as follows: for every data point, we associate a cell z to the pair

¹The inversion may not be always available in closed form, in which case we resort to numerical techniques.

(k_z, x) where k_z is chosen at random from \mathcal{K} and independently from $k_{z'}$ for all $z' \neq z$. This gives us the initial state s_0 .

$$s_0 = z \mapsto (k_z, x) \tag{5}$$

We now need to describe the evolution function Φ . First, assuming that we have a state s and a cell z , we define the following distribution:

$$p_z(k|s) = f(z, T(s)) \tag{6}$$

Assuming that $s(z) = (k, x)$ and that k' is a sample from p_z (hence the name “stochastic” cellular automaton) we define the local update function as:

$$\phi(s, z) = (k', x) \tag{7}$$

$$\text{where } s(z) = (k, x) \text{ and } k' \sim p_z(\cdot|s)$$

That is, the observed data remain unchanged, but we choose a new latent variable according to the distribution p_z induced by the state. We obtain the evolution function of the stochastic cellular automaton by applying the function ϕ uniformly on every cell.

$$\Phi(s) = z \mapsto \phi(s, z) \tag{8}$$

Finally, the SCA algorithm simulates the evolution function Φ starting with s_0 .

As explained earlier, due to our assumption of complete data likelihood belonging to the exponential family, we never have to represent the states explicitly, and instead employ the sufficient statistics.

An implementation can, for example, have two copies of the data structure containing sufficient statistics $T^{(0)}$ and $T^{(1)}$. We do not compute the values $T(\mathbf{z}, \mathbf{x})$ but keep track of the sum as we impute values to the cells/latent variables. During iteration $2t$ of the evolution function, we apply Φ by reading from $T^{(0)}$ and incrementing $T^{(1)}$ as we sample the latent variables (See Figure 1). Then in the next iteration $2t + 1$ we reverse the roles of the data structures, i.e. read from $T^{(1)}$ and increment $T^{(0)}$. We summarize in Algorithm 1.

Algorithm 1 ESCA

- 1: Randomly initialize each cell
 - 2: **for** $t = 0 \rightarrow$ num iterations **do**
 - 3: **for all** cell z **independently in parallel do**
 - 4: Read sufficient statistics from $T^{(t \bmod 2)}$
 - 5: Compute stochastic updates using $p_z(k|s)$
 - 6: Write sufficient statistics to $T^{(t+1 \bmod 2)}$
 - 7: **end for**
 - 8: **end for**
-

Use of such read/write buffers offer a virtually lock-free (assuming atomic increments) implementation scheme for ESCA and is analogous to double-buffering in computer graphics. Although there is a synchronization barrier after each round, its effect is mitigated because each cell’s work depends only upon the sufficient statistics and thus does the same amount of work. There-




Figure 1: Efficient (re)use of buffers

fore, evenly balancing the work load across computational nodes is trivial, even for a heterogeneous cluster.

Furthermore, in the case of discrete latent variable, updating sufficient statistics only requires increments to the data structure $T^{(r)}$ allowing the use of approximate counters [20, 5]. Approximate counters greatly reduce memory costs for the counts: e.g., only 4 or 8 bits per counter. Recent empirical evidence demonstrates that approximate counters preserve statistical performance without compromising runtime performance [32]. In fact, speed often increases because not every increment to the counter results in a write to memory. Note, due to the compression, maintaining two buffers requires less memory than one. Finally, if the latent variables are discrete valued then we can leverage the fast Vose’ alias method [34] to sample. The $O(|\mathcal{K}|)$ construction cost for the alias method can be easily amortized because the rule is homogeneous and thus alias tables can be shared. Details about alias sampling method is provided in Appendix E.

2.4 Wide Applicability of ESCA

As stated previously, ESCA is technically applicable to any model in which the complete data likelihood is in the exponential family. Designing an ESCA algorithm for a model of interest requires simply deriving the S-step for the local update function in the automaton. The S-step is the full conditional (Equation 6) which is easy to derive for many models; for example, mixture models in which (1) the data and parameter components are conjugate and (2) the latent variables and priors are conjugate. We list a few examples of such models in Table 1 and provide additional details in Appendix F. Of course, the extent to which the models enable exploitation of sparsity varies.

ESCA is also applicable to models such as restricted Boltzmann machines (RBMs). For example, if the data were a collection of images, each cell could independently compute the S-step for its respective image. For RBMs the cell would flip a biased coin for each latent variable, and for deep Boltzmann machines, the cells could perform Gibbs sampling. We save a precise derivation and empirical evaluation for future work.

2.5 Understanding the limitations of ESCA

While ESCA has tremendous potential as a computational model for machine learning, in some cases, using it to obtain MAP estimates is not clear.

Consider an Ising model on an d -dimensional torus \mathbb{H}

$$p(x) \propto \prod_{\langle i,j \rangle \in \mathbb{H}} \exp(w_{ij}x_i x_j) \quad (9)$$

in which x_i takes on values in $\{-1, 1\}$. The equilibrium distribution of SCA with a Gibbs update is then [23]

$$q(x) \propto \prod_{\langle i,j \rangle \in \mathbb{H}} \cosh(w_{ij}x_i x_j). \quad (10)$$

Note that the hyperbolic cosine function (\cosh) is symmetric in the sense that $\cosh(r) = \cosh(-r)$. For values $r \geq 0$ \cosh is a good approximation and has a maximum that corresponds to the exponential function; however, for values $r < 0$, the \cosh is a poor approximation for the exponential function.

Let x_1, x_2 be two random variables taking on values in $\{-1, 1\}$. We define a simple two-variable Ising model on a trivial one-dimensional torus:

$$p(x_1, x_2) \propto \exp(x_1 x_2) \quad (11)$$

We can enumerate and quantify the state space under both SCA $q(x_1, x_2)$ and the true distribution $p(x_1, x_2)$:

state	x_1	x_2	$x_1 * x_2$	$q(x_1, x_2) \propto$	$p(x_1, x_2) \propto$
0	-1	-1	1	$\cosh(1)$	$\exp(1)$
1	-1	1	-1	$\cosh(-1)$	$\exp(-1)$
2	1	-1	-1	$\cosh(-1)$	$\exp(-1)$
3	1	1	1	$\cosh(1)$	$\exp(1)$

Since \cosh is symmetric, all states are equally probable for SCA and states 1 and 2 are MAP states. Yet, under the true distribution, they are not. Consequently, SCA with a Gibbs rule for the local evolution function can yield incorrect MAP estimates.

Fortunately, in most cases we are interested in a model over a dataset in which the data is *i.i.d.* That is, we can fix our example as follows. Rather than parallelizing a single Ising model at the granularity of pixels (over a single torus or grid), we instead parallelize the Ising model at the granularity of the data (over multiple tori, one for each image). Then, we could employ Gibbs sampling on each image for the S-step.

2.6 Convergence

We now address the critical question of how the invariant measure of ESCA for the model presented in Section 2.1 is related to the true MAP estimates. First, note that SCA is ergodic [17], a result that immediately applies if we ignore the deterministic components of our automata (corresponding to the observations). Now that we have established ergodicity, we next study the properties of the stationary distribution and find that the modes correspond to MAP estimates.

We make a few mild assumptions about the model:

- The observed data Fisher information is non-singular, i.e. $I(\eta) \succ 0$.
- For the Fisher information for $\mathbf{z}|\mathbf{x}$, we need it to be non-singular and central limit theorem, law of large number to hold, i.e. $\mathbb{E}_{\eta_0}[I_X(\eta_0)] \succ 0$ and

$$\sup_{\eta} \left| \frac{1}{n} \sum_{i=1}^n I_{x_i}(\eta) - \mathbb{E}_{\eta_0}[I_X(\eta)] \right| \rightarrow 0 \text{ as } n \rightarrow \infty$$

- We assume that $\frac{1}{n} \sum_{i=1}^n \nabla_{\eta} \log p(x_i; \eta) = 0$ has at least one solution, let $\hat{\eta}$ be a solution.

These assumptions are reasonable. For example in case of mixture models (or topic models), it just means all component must be exhibited at least once and all components are unique. The details of this case are worked out in Appendix D. Also when the number of parameters grow with the data, e.g., for topic models, the second assumption still holds. In this case, we resort to corresponding result from high dimensional statistics by replacing the law of large numbers with Donsker’s theorem and everything else falls into place.

Consequently, we show ESCA converges weakly to a distribution with mean equal to some root of the score function ($\nabla_{\eta} \log p(x_i; \eta)$) and thus a MAP fixed point by borrowing the results known for SEM [25]. In particular, we have:

Theorem 1 *Let the assumptions stated above hold and $\hat{\eta}$ be the estimate from ESCA. Then as the number of i.i.d. data point goes to infinity, i.e. $n \rightarrow \infty$, we have*

$$\sqrt{n}(\hat{\eta} - \hat{\eta}) \xrightarrow{\mathcal{D}} \mathcal{N}(0, I(\eta_0)^{-1}[I - F(\eta_0)^{-1}]) \quad (12)$$

where $F(\eta_0) = I + \mathbb{E}_{\eta_0}[I_X(\eta_0)](I(\eta_0) + \mathbb{E}_{\eta_0}[I_X(\eta_0)])^{-1}$.

This result implies that SEM flocks around a stationary point under very reasonable assumptions and tremendous computational benefits. Also, for such complicated models, reaching a stationary point is the best that most methods achieve anyway. Now we switch gears to adopt ESCA for LDA and perform some simple experimental evaluations.

3 ESCA for LDA

Topic modeling, and latent Dirichlet allocation (LDA) [2] in particular, have become a must-have of analytics platforms and consequently needs to scale to larger and larger datasets. In LDA, we model each document m of a corpus of M documents as a distribution θ_m that represents a mixture of topics. There are K such topics, and we model each topic k as a distribution ϕ_k over the vocabulary of words that appear in our corpus. Each document m contains N_m words w_{mn} from a vocabulary of size V , and we associate a latent variable z_{mn} to each of the words. The latent variables can take one of K values indicating the topic for the word. Both distributions θ_m and ϕ_k have a Dirichlet prior, parameterized respectively with a constant α and β . See Appendix B for more details.

3.1 Existing systems

Many of the scalable systems for topic modeling are based on one of two core inference methods: the collapsed Gibbs sampler (CGS) [12], and variational inference (VI) [2] and approximations thereof [1]. To scale LDA to large datasets, or for efficiency reasons, we may need to distribute and parallelize them. Both algorithms can be further approximated to meet such implementation requirements.

Collapsed Gibbs Sampling In collapsed Gibbs sampling the full conditional distribution of the latent topic indicators given all the others is

$$p(z_{mn} = k | \mathbf{z}^{-mn}, \mathbf{w}) \propto \frac{(D_{mk} + \alpha) W_{kw_{mn}} + \beta}{T_k + \beta V} \quad (13)$$

where D_{mk} is the number of latent variables in document m that equal k , W_{kv} is the number of latent variables equal to k and whose corresponding word equals v , and T_k is the number of latent variables that equal k , all excluding current z_{mn} .

CGS is a sequential algorithm in which we draw latent variables in turn, and repeat the process for several iterations. The algorithm performs well statistically, and has further benefited from breakthroughs that lead to a reduction of the sampling complexity [37, 16]. This algorithm can be approximated to enable distribution and parallelism, primarily in two ways. One is to partition the data, perform one sampling pass and then assimilate the sampler states, thus yielding an approximate distributed version of CGS (AD-LDA) [24]. Another way is to partition the data and allow each sampler to communicate with a distributed central storage continuously. Here, each sampler sends the differential to the global state-keeper and receives from it the latest global value. A scalable

system built on this principle and leveraging inherent sparsity of LDA is YahooLDA [30]. Further improvement and sampling using alias table was incorporated in lightLDA [39]. Contemporaneously, a nomadic distribution scheme and sampling using Fenwick tree was proposed in F+LDA [38].

Variational Inference In variational inference (VI), we seek to optimize the parameters of an approximate distribution that assumes independence of the latent variables to find a member of the family that is close to the true posterior. Typically, for LDA, document-topic proportions and topic indicators are latent variables and topics are parameter. Then, coordinate ascent alternates between them.

One way to scale VI is stochastic variational inference (SVI) which employs SGD by repeatedly updating the topics via randomly chosen document subsets [13]. Adding a Gibbs step to SVI introduces sparsity for additional efficiency [19]. In some ways this is analogous to our S-step, but in the context of variational inference, the conditional is much more expensive to compute, requiring several rounds of sampling.

Another approach, CVB0, achieves scalability by approximating the collapsed posterior [31]. Here, they minimize the free energy of the approximate distribution for a given parameter γ_{mnk} and then use the zero-order Taylor expansion [1].

$$\gamma_{mnk} \propto (D_{mk} + \alpha) \times \frac{W_{kw_{mn}} + \beta}{T_k + \beta V} \quad (14)$$

where D_{mk} is the fractional contribution of latent variables in document m for topic k , W_{kv} is the contribution of latent variables for topic k and whose corresponding word equals v , and T_k is the the contribution of latent variables for topic k . Inference updates the variational parameters until convergence. It is possible to distribute and parallelize CVB0 over tokens [1]. VI and CVB0 are the core algorithms behind several scalable topic modeling systems including Mr.LDA [40] and the Apache Spark machine-learning suite.

Remark It is worth noticing that Gibbs sampling and variational inference, despite being justified very differently, have at their core the very same formulas (shown in a box in formula (13) and (14)). Each of which are literally deciding how important is some topic k to the word v appearing in document m by asking the questions: “How many times does topic k occur in document m ?”, “How many times is word v associated with topic k ?”, and “How prominent is topic k overall?”. It is reassuring that behind all the beautiful mathematics, something simple and intuitive is happening. As we see next, ESCA addresses the same questions via analogous formulas for SEM.

3.2 An ESCA Algorithm for LDA

To re-iterate, the point of using such a method for LDA is that the parallel update dynamics of the ESCA gives us an algorithm that is simple to parallelize, distribute and scale. In the next section, we will evaluate how it works in practice. For now, let us explain how we design our SCA to analyze data.

We begin by writing the stochastic EM steps for LDA (derivation is in Appendix B):

E-step: independently in parallel compute the conditional distribution locally:

$$q_{mnk} = \frac{\theta_{mk} \phi_{kw_{mn}}}{\sum_{k'=1}^K \theta_{mk'} \phi_{k'w_{mn}}} \quad (15)$$

S-step: independently in parallel draw z_{ij} from the categorical distribution:

$$z_{mn} \sim \text{Categorical}(q_{mn1}, \dots, q_{mnK}) \quad (16)$$

M-step: independently in parallel compute the new parameter estimates:

$$\begin{aligned} \theta_{mk} &= \frac{D_{mk} + \alpha - 1}{N_m + K\alpha - K} \\ \phi_{kv} &= \frac{W_{kv} + \beta - 1}{T_k + V\beta - V} \end{aligned} \quad (17)$$

We simulate these inference steps in ESCA, which is a dynamical system with evolution function $\Phi : \mathcal{S} \rightarrow \mathcal{S}$ over the state space \mathcal{S} . For LDA, the state space \mathcal{S} is

$$\mathcal{S} = \mathcal{Z} \rightarrow \mathcal{K} \times \mathcal{M} \times \mathcal{V} \quad (18)$$

where \mathcal{Z} is the set of cell identifiers (one per token in our corpus), \mathcal{K} is a set of K topics, \mathcal{M} is a set of M document identifiers, and \mathcal{V} is a set of V identifiers for the vocabulary words.

The initial state s_0 is the map defined as follows: for every occurrence of the word v in document m , we associate a cell z to the triple (k_z, m, v) where k_z is chosen uniformly at random from \mathcal{K} and independently from $k_{z'}$ for all $z' \neq z$. This gives us

$$s_0 = z \mapsto (k_z, m, v) \quad (19)$$

We now need to describe the evolution function Φ . First, assuming that we have a state s and a cell z , we define the following distribution:

$$p_z(k|s) \propto (D_{mk} + \alpha) \times \frac{W_{kv} + \beta}{T_k + \beta V} \quad (20)$$

where $D_{mk} = \left| \{ z \mid \exists v. s(z) = (k, m, v) \} \right|$, $W_{kv} = \left| \{ z \mid \exists m. s(z) = (k, m, v) \} \right|$, and $T_k = \left| \{ z \mid \exists m. \exists v. s(z) = (k, m, v) \} \right|$. Note that we have chosen our local update rule slightly different without an offset of -1 for the counts corresponding to

the mode of the Dirichlet distributions and requiring $\alpha, \beta > 1$. Instead, our local update rule allows us to have the relaxed requirement $\alpha, \beta > 0$ which is more common for LDA inference algorithms.

Assuming that $s(z) = (k, m, v)$ and that k' is a sample from p_z (hence the name “stochastic” cellular automaton) we define the local update function as:

$$\phi(s, z) = (k', m, v) \quad (21)$$

where $s(z) = (k, m, v)$ and $k' \sim p_z(\cdot | s)$

That is, the document and word of the cell remain unchanged, but we choose a new topic according to the distribution p_z induced by the state. We obtain the evolution function of the stochastic cellular automaton by applying the function ϕ uniformly on every cell.

$$\Phi(s) = z \mapsto \phi(s, z) \quad (22)$$

Finally, the SCA algorithm simulates the evolution function Φ starting with s_0 . Of course, since LDA’s complete data likelihood is in the exponential family, we never have to represent the states explicitly, and instead employ the sufficient statistics.

Our implementation has two copies of the count matrices D^i , W^i , and T^i for $i = 0$ or 1 (as in CGS or CVB0, we do not compute the values D_{ik} , W_{kv} , and T_k but keep track of the counts as we assign topics to the cells/latent variables). During iteration i of the evolution function, we apply Φ by reading $D^{i \bmod 2}$, $W^{i \bmod 2}$, and $T^{i \bmod 2}$ and incrementing $D^{(i+1) \bmod 2}$, $W^{(i+1) \bmod 2}$, and $T^{(i+1) \bmod 2}$ as we assign topics.

3.3 Advantages of ESCA for LDA

The positive consequences of ESCA as a choice for inference on LDA are many:

- Our memory footprint is minimal since we only store the data and sufficient statistics. In contrast to MCMC methods, we do not store the assignments to latent variables \mathbf{z} . In contrast to variational methods, we do not store the variational parameters γ . Further, variational methods require K memory accesses (one for each topic) per word. In contrast, the S-step ensures we only have a single access (for the sampled topic) per word. Such reduced pressure on the memory bandwidth can improve performance significantly for highly parallel applications.
- We can further reduce the memory footprint by compressing the sufficient statistics with approximate counters [20, 5]. This is possible because updating the sufficient statistics only requires increments as in Mean-for-Mode [32]. In contrast, CGS decrements counts, preventing the use of approximate counters.
- Our implementation is lock-free (in that it does not use locks, but assumes atomic increments) because the double buffering ensures we never read or write

to the same data structures. There is less synchronization, which at scale is significant.

- Finally, our algorithm is able to fully benefit from Vose’s alias method [34] because homogeneous update rule for SCA ensures that the cost for constructing the alias table is amortized across the cells. To elaborate, the SCA update Equation (20) decomposes as

$$p_z(k|s) \propto \left[D_{mk} \frac{W_{kv} + \beta}{T_k + \beta V} \right] + \left[\alpha \frac{W_{kv} + \beta}{T_k + \beta V} \right] \quad (23)$$

allowing us to treat it as a discrete mixture and divide the sampling procedure into a two steps. First, we toss a biased coin to decide which term of the equation to sample, and second, we employ a specialized sampler depending on the chosen term. The first term is extremely sparse (documents comprise only a small handful of topics) and a basic sampling procedure suffices. The second term is not sparse, but is independent of the current document m and depends only on the W and T matrices. Moreover, as mentioned earlier, during iteration i , we will be only reading values from non-changing $W^{i \bmod 2}$, and $T^{i \bmod 2}$ matrices. As a result, at the start of each iteration we can precompute, from the W and T matrices, tables for use with Vose’s alias method, which enables sampling from the second term in a mere 3 CPU operations. Thus, the evolution for ESCA is extremely efficient.

3.3.1 Connection to SGD

We can view ESCA as implicit SGD on MAP for LDA. This connection alludes to the convergence rate of ESCA. To illustrate, we consider θ only. As pointed out in [36, 29], one EM step is:

$$\theta_m^+ = \theta_m + M \frac{\partial \log p}{\partial \theta_{mk}}$$

which is gradient descent with a Frank-Wolfe update and line search. Similarly, for ESCA using stochastic EM, one step is

$$\theta_{mk}^+ = \frac{D_{mk}}{N_m} = \frac{1}{N_m} \sum_{n=1}^{N_m} \delta(z_{mn} = k)$$

Again vectorizing and re-writing as earlier:

$$\theta_m^+ = \theta_m + M g$$

where $M = \frac{1}{N_m} [\text{diag}(\theta_m) - \theta_m \theta_m^T]$ and $g = \frac{1}{\theta_{mk}} \sum_{n=1}^{N_m} \delta(z_{mn} = k)$. The vector g can be shown to be an unbiased noisy estimate of the gradient, i.e.

$$\mathbb{E}[g] = \frac{1}{\theta_{mk}} \sum_{n=1}^{N_m} \mathbb{E}[\delta(z_{ij} = k)] = \frac{\partial \log p}{\partial \theta_{mk}}$$

Thus, a single step of SEM on our SCA is equivalent to a single step of SGD. Consequently, we could further embrace the connection to SGD and use a subset of the

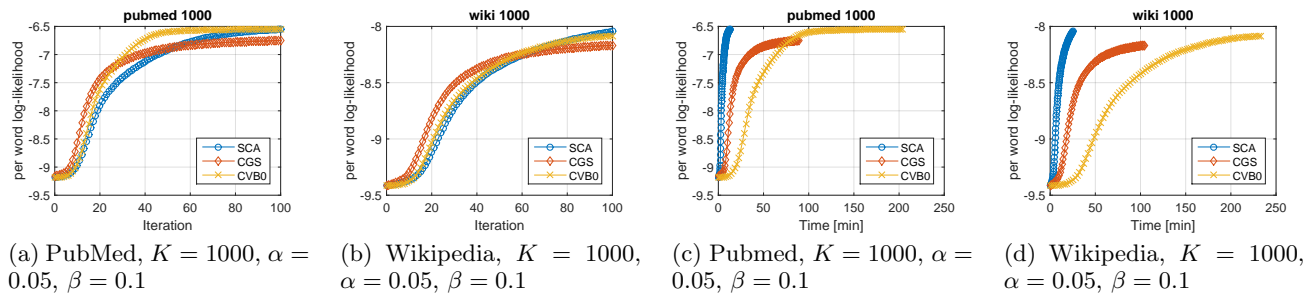


Figure 2: Evolution of log likelihood on Wikipedia and Pubmed over number of iterations and time.

data for the S and M steps, similar to incremental EM [22]. Note that in the limit in which batches comprise just a single token, the algorithm emulates a collapsed Gibbs sampler. This interpretation strengthens the theoretical justification for many existing approximate Gibbs sampling approaches.

4 Experiments

To evaluate the strength and weaknesses of our algorithm, we compare against parallel and distributed implementations of CGS and CVB0. We also compare our results to performance numbers reported in the literature including those of F+LDA and lightLDA.

Software & hardware All three algorithms are implemented in simple C++11. We implement multithreaded parallelization within a node using the work-stealing Fork/Join framework, and the distribution across multiple nodes using the process binding to a socket over MPI. We also implemented a version of ESCA with a sparse representation for the array D of counts of topics per documents and Vose’s alias method to draw from discrete distributions. We run our experiments on a small cluster of 8 Amazon EC2 c4.8xlarge nodes connected through 10Gb/s Ethernet. Each node has a 36 virtual threads per node. For random number generation we employ Intel® Digital Random Number Generators through instruction RDRAND, which uses thermal noise within the silicon to output a random stream of bits at 3 Gbit/s, producing true random numbers.

Datasets We experiment on two public datasets, both of which are cleaned by removing stop words and rare words: PubMed abstracts and English Wikipedia. To demonstrate scalability, run on 100 copies of English Wikipedia and a third proprietary dataset.

Dataset	V	M	Tokens
PubMed	141,043	8,200,000	737,869,085
Wikipedia	210,233	6,631,176	1,133,050,514
Large	~140,000	~3 billion	~171 billion

Evaluation To evaluate the proposed method we use predicting power as a metric by calculating the per-word log-likelihood (equivalent to negative log of perplexity) on 10,000 held-out documents conditioned on the trained model. We set $K = 1000$ to demonstrate performance for a large number of topics. The hyper parameters are set as $\alpha = 50/K$ and $\beta = 0.1$ as suggested in [12]; other systems such as YahooLDA and Mallet also use this as the default parameter setting. The results are presented in Figure 2 and some more experiments in Appendix G.

Finally, for the larger datasets, our implementation of ESCA (only 300 lines of C++) processes more than **1 billion tokens per second** (tps) by using a 20-node cluster. In comparison, some of the best existing systems achieve 112 million tps (F+LDA, personal communication) and 60 million tps (lightLDA) [39] using more hardware. Detailed Appendix G Table 3)

5 Discussion

We have described a novel inference method for latent variable models that simulates a stochastic cellular automaton. The equilibrium of the dynamics are MAP fixed points and the algorithm has many desirable computational properties: it is embarrassingly parallel, memory efficient, and like HOGWILD! [27], is virtually lock-free. Further, for many models, it enables the use of approximate counters and the alias method. Thus, we were able to achieve an order of magnitude speed-up over the current state-of-the-art inference algorithms for LDA with accuracy comparable to collapsed Gibbs sampling.

In general, we cannot always guarantee the correct invariant measure [7], and found that parallelizing improperly causes convergence to incorrect MAP fixed points. Even so, SCA is used for simulating Ising models in statistical physics [33]. Interestingly, in previous work [15], it has been shown that stochastic cellular automata are closely related to equilibrium statistical models and the stationary distribution is known for a large class of finite stochastic cellular automata.

References

- [1] Arthur Asuncion, Max Welling, Padhraic Smyth, and Yee Whye Teh. On smoothing and inference for topic models. In *Proc. Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI '09, pages 27–34, Arlington, Virginia, USA, 2009. AUAI Press.
- [2] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, March 2003.
- [3] J. Canny. Gap: a factor model for discrete data. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 122–129. ACM, 2004.
- [4] Gilles Celeux and Jean Diebolt. The sem algorithm: a probabilistic teacher algorithm derived from the em algorithm for the mixture problem. *Computational statistics quarterly*, 2(1):73–82, 1985.
- [5] Miklós Csűrös. Approximate counting with a floating-point counter. In M. T. Thai and Sartaj Sahni, editors, *Computing and Combinatorics (COCOON 2010)*, number 6196 in Lecture Notes in Computer Science, pages 358–367. Springer Berlin Heidelberg, 2010. See also <http://arxiv.org/pdf/0904.3062.pdf>.
- [6] Rajarshi Das, Manzil Zaheer, and Chris Dyer. Gaussian lda for topic models with word embeddings. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 795–804, Beijing, China, July 2015. Association for Computational Linguistics.
- [7] Donald A. Dawson. Synchronous and asynchronous reversible Markov systems. *Canadian mathematical bulletin*, 17:633–649, 1974.
- [8] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, January 2008.
- [9] Anton K Formann and Thomas Kohlmann. Latent class analysis in medical research. *Statistical methods in medical research*, 5(2):179–211, 1996.
- [10] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1995.
- [11] Joseph Gonzalez, Yucheng Low, Arthur Gretton, and Carlos Guestrin. Parallel gibbs sampling: from colored fields to thin junction trees. In *International Conference on Artificial Intelligence and Statistics*, pages 324–332, 2011.
- [12] T.L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101:5228–5235, 2004.
- [13] Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, May 2013.
- [14] Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Mach. Learn.*, 37(2):183–233, November 1999.
- [15] Joel L. Lebowitz, Christian Maes, and Eugene R. Speer. Statistical mechanics of probabilistic cellular automata. *Journal of statistical physics*, 59:117–170, April 1990.
- [16] Aaron Q. Li, Amr Ahmed, Sujith Ravi, and Alexander J. Smola. Reducing the sampling complexity of topic models. In *20th ACM SIGKDD Intl. Conf. Knowledge Discovery and Data Mining*, 2014.
- [17] Pierre-Yves Louis. *Automates Cellulaires Probabilistes : mesures stationnaires, mesures de Gibbs associées et ergodicité*. PhD thesis, Université des Sciences et Technologies de Lille and il Politecnico di Milano, September 2002.
- [18] Jean Mairesse and Irène Marcovici. Around probabilistic cellular automata. *Theoretical Computer Science*, 559:42–72, November 2014.
- [19] David Mimno, Matt Hoffman, and David Blei. Sparse stochastic inference for latent dirichlet allocation. In John Langford and Joelle Pineau, editors, *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, ICML '12, pages 1599–1606, New York, NY, USA, July 2012. Omnipress.
- [20] Robert Morris. Counting large numbers of events in small registers. *Commun. ACM*, 21(10):840–842, October 1978.
- [21] R. Neal. Markov chain sampling methods for dirichlet process mixture models. Technical Report 9815, University of Toronto, 1998.
- [22] Radford M Neal and Geoffrey E Hinton. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pages 355–368. Springer, 1998.

- [23] A. U. Neumann and B. Derrida. Finite size scaling study of dynamical phase transitions in two dimensional models: Ferromagnet, symmetric and non symmetric spin glasses. *J. Phys. France*, 49:1647–1656, 08 1988.
- [24] David Newman, Arthur Asuncion, Padhraic Smyth, and Max Welling. Distributed algorithms for topic models. *J. Machine Learning Research*, 10:1801–1828, December 2009. <http://dl.acm.org/citation.cfm?id=1577069.1755845>.
- [25] Søren Feodor Nielsen. The stochastic em algorithm: estimation and asymptotic results. *Bernoulli*, pages 457–489, 2000.
- [26] Sam Patterson and Yee Whye Teh. Stochastic gradient riemannian langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, pages 3102–3110, 2013.
- [27] B. Recht, C. Re, S.J. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In Peter Bartlett, Fernando Pereira, Richard Zemel, John Shawe-Taylor, and Kilian Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 693–701, 2011.
- [28] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 09 1951.
- [29] Ruslan Salakhutdinov, Sam Roweis, and Zoubin Ghahramani. Relationship between gradient and em steps in latent variable models.
- [30] Alexander Smola and Shравan Narayanamurthy. An architecture for parallel topic models. *Proc. VLDB Endowment*, 3(1-2):703–710, September 2010.
- [31] Whye Yee Teh, David Newman, and Max Welling. A collapsed variational Bayesian inference algorithm for latent Dirichlet allocation. In *Advances in Neural Information Processing Systems 19*, NIPS 2006, pages 1353–1360. MIT Press, 2007.
- [32] Jean-Baptiste Tristan, Joseph Tassarotti, and Guy L. Steele Jr. Efficient training of LDA on a GPU by Mean-For-Mode Gibbs sampling. In *32nd International Conference on Machine Learning*, volume 37 of *ICML 2015*, 2015. Volume 37 of the Journal in Machine Learning Research: Workshop and Conference Proceedings.
- [33] Gérard Y. Vichniac. Simulating physics with cellular automata. *Physica D: Nonlinear Phenomena*, 10(1-2):96–116, January 1984.
- [34] Michael D Vose. A linear algorithm for generating random numbers with a given distribution. *Software Engineering, IEEE Transactions on*, 17(9):972–975, 1991.
- [35] Max A Woodbury, Jonathan Clive, and Arthur Garson. Mathematical typology: a grade of membership technique for obtaining disease definition. *Computers and biomedical research*, 11(3):277–298, 1978.
- [36] Lei Xu and Michael I Jordan. On convergence properties of the em algorithm for gaussian mixtures. *Neural computation*, 8(1):129–151, 1996.
- [37] Limin Yao, David Mimno, and Andrew McCallum. Efficient methods for topic model inference on streaming document collections. In *Proc. 15th ACM SIGKDD Intl. Conf. Knowledge Discovery and Data Mining*, KDD '09, pages 937–946, New York, 2009. ACM.
- [38] Hsiang-Fu Yu, Cho-Jui Hsieh, Hyokun Yun, SVN Vishwanathan, and Inderjit S Dhillon. A scalable asynchronous distributed algorithm for topic modeling. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1340–1350. International World Wide Web Conferences Steering Committee, 2015.
- [39] Jinhui Yuan, Fei Gao, Qirong Ho, Wei Dai, Jintian Wei, Xun Zheng, Eric Po Xing, Tie-Yan Liu, and Wei-Ying Ma. Lightlda: Big topic models on modest computer clusters. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1351–1361. International World Wide Web Conferences Steering Committee, 2015.
- [40] Ke Zhai, Jordan Boyd-Graber, Nima Asadi, and Mohamad L Alkhouja. Mr. lda: A flexible large scale topic modeling package using variational inference in mapreduce. In *Proceedings of the 21st international conference on World Wide Web*, pages 879–888. ACM, 2012.