

Parallel LDA Through Optimized Synchronous Communication Methods

Bingjing Zhang, Bo Peng, Judy Qiu
 Computer Science Department
 Indiana University
 Bloomington, IN, USA
 zhangbj, pengb, xqiu@indiana.edu

Abstract—Selecting an appropriate distributed processing framework can be difficult for developers building large-scale machine learning applications. That is because all these tools provide various kinds of parallelism patterns and suggest different communication strategies to synchronize local and global model data distributed among parallel nodes. There is no clear answer to determine which strategy might be suitable depending on the data and application. Taking Latent Dirichlet Allocation (LDA) as an example, contemporary implementations often choose asynchronous communication methods to synchronize the model data. However, our observations show that the asynchronous communication still has very high overhead and the characteristics of the LDA training datasets encourage us to use optimized synchronous collective communication methods instead. The results show that with data parallelism only, our "lda-lgs" implementation can be (%) faster compared to Yahoo! LDA. With model parallelism, our "lda-rtt" implementation has similar speed compared with Petuum LDA on a uni-gram model with 1 million words and 10k topics but (%) faster on a bi-gram model with 20 million words and 500 topics.

I. INTRODUCTION

One challenge of parallel machine learning applications is that while training data can be split into parallel workers, the model data that all local computations depend on is growing progressively and generates significant synchronization overhead. Currently two types of parallelism are used to solve this problem (see Fig. 1a):

Data Parallelism The global model is distributed on a set of servers or on existing parallel workers. Each worker samples on a local model and updates it through the synchronization between local models and the global model.

Model Parallelism In addition to using data parallelism, the global model data is split between parallel workers and rotated during the sampling.

In Latent Dirichlet Allocation (LDA) [1], the model synchronization is important because a faster communication method not only reduces the resulting overhead, but also speeds up the model convergence rate, shrinks the model size, and shortens the computation time in later iterations. Though both synchronous and asynchronous methods (see Fig. 1b) can cause the model to converge without affecting the correctness of the algorithm, it is unclear which strategy performs better for LDA applications. Asynchronous communication is popular because it avoids the overhead of global waiting between parallel workers and that of local waiting between

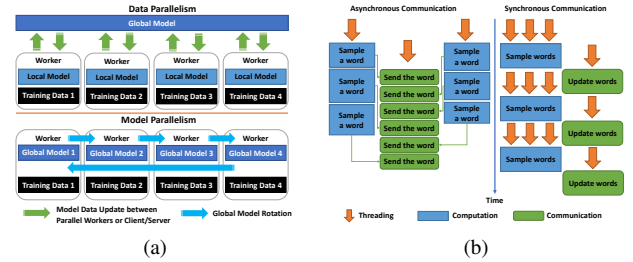


Fig. 1. (a) Data Parallelism vs. Model Parallelism and (b) Asynchronous Communication vs. Synchronous Communication in LDA

computation threads and communication threads. In data parallelism, asynchronous communication allows local computation to continue without waiting for the completion of updating the global model from all parallel workers per iteration. In model parallelism, though model rotation is synchronous, per word sampling and sending can still overlap without waiting on each worker, demonstrating asynchronous communication.

However, after studying the characteristics of LDA data, we have identified that the counts of each word in the training documents fall under the power-law distribution. As a result, when data parallelism is used, many words in the global model will display on all the workers' local models, and this generates "one-to-all" communication patterns during the synchronization. Similarly, in model parallelism, as the size of the global model data expands, each worker needs to handle more data transference. These observations inspired us to apply routing optimized synchronous communication operations to improve the the LDA model update speed.

Our synchronous communication methods utilize the model data distribution characteristics and routing optimization in conjunction. Furthermore, we overlapped the computation and communication steps to reduce the overhead of the global/local waiting. These ideas are implemented in Harp [2], a collective communication library on Hadoop. Harp has already integrated several collective communication patterns from different parallel processing frameworks in a unified abstraction. However, all the current patterns cannot abstract either the local/global model synchronization in data parallelism or the model rotation in model parallelism. As such, we abstracted three other communication patterns called "syncLocalWithGlobal",

“syncGlobalWithLocal”, and “rotateGlobal” in which our new ideas are embedded. The new patterns are very generalizable so that they can be applied not only to LDA applications but also to many other machine learning applications. We implemented one LDA application which uses “syncLocalWithGlobal” and “syncGlobalWithLocal” to perform data parallelism and another which uses “rotateGlobal” to perform model parallelism. We compared our implementations with other implementations based on asynchronous communication methods, such as Yahoo! LDA [3] and Petuum LDA [4], on several datasets. The results show that optimized synchronous communication methods can reduce communication overhead and improve model convergence speed.

The following sections describe: the cost model of LDA algorithm (Section 2), the synchronous communication methods (Section 3), the implementation of Harp-LDA (Section 4), the performance results of our implementation (Section 5), the related work on parallel LDA (Section 6), and our conclusions (Section 7).

II. COST MODEL

A. LDA model

LDA is a generative probabilistic data modeling technique. Training data are abstracted as a document collection where each document is a bag of words. LDA models the data by introducing latent topics, which tries to capture the underlining semantic connections and structures inside the data. In LDA model, a document is a mixture of latent topics, and each topic is a multinomial distribution over words. In the generative process, for document j , we first draw a topic distribution θ_j from a Dirichlet with parameter α . Then for each word i in this document, we draw a topic $z_{ij} = k$ from the multinomial distribution with parameter θ_j . Finally, word x_{ij} is drawn from a multinomial $\phi_{wk|k=z_{ij}}$, which also derives from a Dirichlet with parameter β . Here, the words x_{ij} are observed variables, θ, ϕ, z are latent variables, and α and β are hyper parameters.

The purpose of LDA inference is to compute the posterior distribution of the latent variables given the observed variables. There are many approximate inference algorithms. In a practice on large data, Collapsed Gibbs Sampling (CGS) [5] displays high scalability. Collapse is a procedure used to integrate out θ, ϕ and sample only the latent variables z . Gibbs Sampling is one kind of Markov Chain Monte Carlo algorithm for inference. There are three phases, the initialization, burn-in, and stationary phase.

In initialization, each word is initialized by a random topic denoted as z_{ij} . Then it begins to reassign topics to each word w_{ij} according to the conditional probability of z_{ij} , which is henceforth called sampling.

$$p(z_{ij} = k | z^{-ij}, x, \alpha, \beta) \propto \frac{N_{wk}^{-ij} + \beta}{\sum_w N_{wk}^{-ij} + V\beta} (N_{kj}^{-ij} + \alpha) \quad (1)$$

Here, superscript $-ij$ means that the corresponding word is excluded in the counts. V is vocabulary size. N_{wk} is the count

of word w assigned to topic k , and N_{kj} is the count of topic k assigned in document j , which are sufficient statistics for the latent variable θ and ϕ . The latent variables can be represented by three matrices Z_{ij} , N_{wk} and N_{kj} , which are model data. Intuitively, by equation(1), with higher probability a word will be assigned to the topic that has been assigned to its co-occurring words. Therefore, sampling by the latest model data of co-occurring words is critical for convergence, and that is why synchronization is so important in a parallel LDA trainer.

Hyper parameters α and β are also called concentration parameters, which control the topic density in the final model. The larger the α and β , the more topics can be drawn into a document and assigned to a word, and the more non-zero cells in each row of the N_{wk} and N_{kj} matrices. Although a useful LDA trainer often has the feature of α and β optimization dynamically tuned to fit the training data, in this paper, we skip such a feature and use symmetric α and β both fixed to a common used value 0.01 to exclude the complex effects on performance caused by their dynamics.

Latent variables will gradually converge in the process of iterative sampling. This is the phase where burn-in occurs and finally reaches the stationary state. From that point, we can draw samples from the sampling process and use them to calculate the posterior distribution.

To evaluate the quality of the final model learned by LDA, held-out testsets are often used, taking likelihood or perplexity as the accuracy metrics. In this paper, we only use the model data likelihood on the training dataset to monitor the convergence of the LDA trainer, which is consistent with the held-out testset results in our experiments, only much faster.

Sampling on z_{ij} in CGS is a strictly sequential process. AD-LDA [6] is the seminal work allowing us to relax this sequential sampling requirement. It assumes that the dependence between one topic assignment z_{ij} and another z_{ij} is weak in that different words in different documents are sampled concurrently. In AD-LDA, training data are partitioned into n subsets, with n Gibbs Samplers running parallel on each collection, and each sampler synchronizing its model data with others at certain time points. This parallel version produces a useful model, establishing the foundation of large-scale parallel CGS implementations of LDA trainers on large-scale data in practice.

B. Performance Factors

Many factors are related to the performance of a LDA trainer.

Sampling Algorithm Computation complexity of a sampling algorithm basically determines the overall performance. Although there is a $\mathcal{O}(1)$ sampling algorithm, LightLDA [7], proposed in the literature, we focus on SparseLDA [8], which is an optimized CGS sampling algorithm mostly used in the state-of-the-art LDA trainers, in order to make a broader

comparison. SparseLDA splits the equation (1) into three parts:

$$p(z_{ij} = k | rest) \propto \frac{N_{wk}^{-ij}(N_{kj}^{-ij} + \alpha) + \beta * N_{kj}^{-ij} + \alpha\beta}{\sum_w N_{wk}^{-ij} + V\beta} \quad (2)$$

The denominator is a constant when sampling on one word. The third part of the numerator is also a constant; the second part is non-zero only when N_{kj} is non-zero, and the first part is non-zero only when N_{wk} is non-zero. In naive CGS sampling, the conditional probability will compute K times, while in SparseLDA, the computation can be decreased to non-zero items number in N_{wk} and N_{kj} , which are much smaller than K on average.

We found that in practice, the sampling performance is more memory bounded than computation bounded, for the computation is very simple and memory access to two large matrices is not by its nature cache friendly. Furthermore, CGS has a feature of exchangeability that permits the order of word sampling to be changed. In practice, sampling can take the order by row or column on the document-word matrix. Equation(2) is the form optimized for row order, called sample-by-doc. In this case, N_{kj} can be cached for the words in the same row, and the computation complexity in terms of amortized random memory access time is $\mathcal{O}(\sum_k \mathbb{1}(N_{wk} \neq 0))$. Symmetrically, sample-by-word will have the complexity of $\mathcal{O}(\sum_k \mathbb{1}(N_{kj} \neq 0))$.

Parallelism Strategy Data partition on the training data, which is a document-word matrix, can be done either in the rows or the columns. If data are partitioned by rows, each subset data has its local z , N_{kj} , N_{wk} model data and only N_{wk} needs to be synchronized with others. In general applications, the row number is much larger than the column number, so partition by rows will generate a smaller model data size. We only refer to the shared word-topic matrix as model data.

There are many possible communication strategies which control how to do model data synchronization between parallel workers. Modern clusters allow two levels of parallelism: inter-node and inner-node parallelism. In this paper, we focus on inter-node parallelism by exploring the differences between the communication strategies.

Cluster configurations include nodes number N and networking bandwidth B , memory size M for each node, and thread number T for each node. As many-core technology brings forth more powerful machines to bear complicated computation applications, large-scale machine learning applications will benefit. Relatively small numbers of N with a large number of T can reach high scale parallelism, which is more like a traditional HPC cluster than a cloud cluster.

Data Property Training data can be characterized by the total numbers of tokens, denoted as W , and the number of documents, denoted as D . The model data N_{wk} is a $V * K$ matrix and N_{kj} is a $D * K$ matrix, where V is the vocabulary size and K is the topic number.

LDA is an iterative algorithm. It keeps sampling on the training data and updating (synchronizing) the model data until it converges. In computation, one iteration is one pass on

sampling the training data. In synchronization, one iteration is one pass to synchronize all the model data. As we described above, both parts are highly related to the model data size, not in terms of the matrix dimension but the non-zero items count.

C. Model Data

Model Size Power law distribution is a general phenomenon. It has another equal form for text data as Zipf's law, where the frequency of a word is proportional to the reciprocal of its rank.

$$freq(i) = C * i^{-\lambda} \quad (3)$$

Here, i is word rank, and λ is near 1.

There are a total of V unique words in the training data. We then have:

$$\begin{aligned} W &= \sum_{i=1}^V (freq(i)) = \sum_{i=1}^V (C * i^{-\lambda}) \\ &\approx C * (\ln(V) + \gamma + \frac{1}{2V}) \end{aligned} \quad (4)$$

If λ is 1, this is the partial sum of harmonic series which have logarithmic growth, where γ is the EulerMascheroni constant ≈ 0.57721 .

Model data, $V * K$, is a very large but sparse matrix. In a general setting, V is 1M, K is 1K, while for big models it can even reach 1M*1M. The non-zero cell count of the matrix is the true model size, denoted as S , $S << V * K$.

In the initialization of CGS, word-topic count matrix is initialized by random topic assignment for each work. So the word i will get $\max(K, freq(i))$ non-zero cells. If $freq(J) = K$, $J = C/K$, we get:

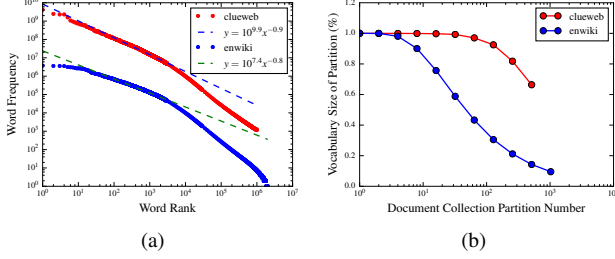
$$\begin{aligned} S_{init} &= \sum_{i=1}^J K + \sum_{i=J+1}^V freq(i) = W - \sum_{i=1}^J freq(i) + \sum_{i=1}^J K \\ &= C * (\ln V + \ln K - \ln C + 1) \end{aligned} \quad (5)$$

The true model size S_{init} is logarithmic to matrix size $V * K$. This does not mean S_{init} is small, for the constant $C = freq(1)$ can be very large; even $C * \ln(V * K)$ can be huge. An increase of dimension in the model will not increase the model data size dramatically.

With the progress of iterations and algorithm convergence, the model data size will shrink. The concentration parameters α and β control the final sparsity of the topic distribution. When a stationary state is reached, the average count value will drop to a certain small constant ratio of K , with the constant δ determined by the properties of the training data itself.

$$S_{final} = \text{mean}(\text{word} - \text{topiccount}) * V = \delta * K * V \quad (6)$$

Model Data Partition After training data is partitioned to each node of the cluster, a local model data S' will be built up and used in local computation. This local model data should synchronize with global model data S frequently to make the training process converge. In fact, the synchronization frequency is highly relevant to the final model accuracy.



First, “syncGlobalWithLocal” uses the data in local tables to synchronize the data in global tables. This operation will reduce the partitions from local tables to the global table. Secondly, “syncLocalWithGlobal” uses the data in global tables to synchronize local tables. Based on the needs of partitions in local tables, this operation will redistribute the partitions in the global table to local tables. If one partition is required by all the workers, it will be broadcasted.

Lastly, “rotateGlobal” will consider workers in a ring topology and shift the partitions in the global table owned by one worker to the right neighbor worker and then receive the partitions from the left neighbor. When the operation is completed, the contents of the distributed dataset in the global tables won’t change, but each worker will hold a different set of partitions. Since each worker only talks to its neighbors, “rotateGlobal” can transmit global data in parallel without any network conflicts.

B. The Applicability of Synchronous Communication Methods

“syncGlobalWithLocal” and “syncLocalWithGlobal” are abstracted from data parallelism, and “rotateGlobal” is abstracted from model parallelism. However, these operations are not limited to the communication patterns in parallel LDA. Instead, they can be applied to many other machine learning algorithms with big model data.

A matrix can be drawn to describe each worker’s requirements on the global model data in the parallel computation per iteration. In this matrix, each row represents a worker, each column represents a global data partition, and each element shows the requirements of the partition in the local computation. Based on the density of this computation relation matrix, we can choose proper operations in different applications. If the matrix is dense, we suggest using the “rotateGlobal” operation. Using k-means clustering as an example, the global model data are the centroids, and the local computation needs all the centroids data. Thus “rotateGlobal” allows each worker to access all the centroids data efficiently. If the matrix is sparse, using “syncGlobalWithLocal” and “syncLocalWithGlobal” is a superior solution. For example, in graph algorithms such as PageRank, the global model data are the vertices’ page-rank values and counts of out-edges. The local computation goes through each edge and calculates the partial result of the new page-rank values. Then “syncGlobalWithLocal” can be used to update global page-rank values. In the next iteration, we can use “syncLocalWithGlobal” to fetch the new global page-rank values to each local computation.

IV. HARP-LDA IMPLEMENTATION

A. Partition Training Data And Initialize Model Data

For the training data, we split the documents into files evenly. For the model data, since words with high frequency can dominate the computation and communication, we partition the global model based on the frequency of words in the training dataset. During the preprocessing of the training data, each word is given an ID based on their frequency starting

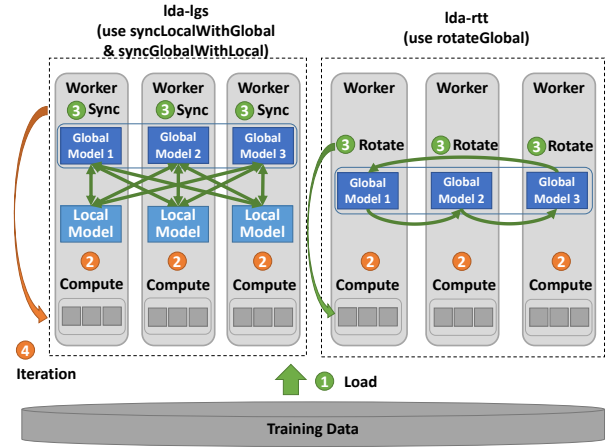


Fig. 3. Internode Parallelism (data loading (step 1) and iteration (step 4) are common procedures for both implementations)

from 0. The lower the occurrence of the word, the higher the ID. Then we partition the words’ topic counts using range-based partitioning. Assuming each partition contains m words, Partition 0 contains words with IDs from 0 to $m - 1$, and Partition 1 contains words with IDs from m to $2m - 1$, and so on. As a result, the partitions with low IDs contain the words with the highest frequency. The initial global model is generated by randomly assigning each token to a topic and aggregated through “syncGlobalWithLocal”. The mapping between partition IDs and worker IDs is calculated based on the modulo operation. Assuming there is a worker with ID w among a total of N workers, the partitions contained on this worker are Partition w , Partition $w + N$, Partition $w + 2N$, and so on. In this way, each worker contains a number of words whose frequencies rank from high to low.

B. Inter-node Parallelism

During iterations of the sampling, we use two different approaches to update the global model which results in two implementations (See Fig. 3). One implementation, named “lda-lgs”, follows data parallelism and uses “syncGlobalWithLocal” paired with “syncLocalWithGlobal” operations. The other implementation, named “lda-rtt”, follows model parallelism and uses “rotateGlobal” operation.

During the sampling of “lda-lgs”, each worker updates the local model and tracks the difference generated in another table. Once the sampling is done, “syncGlobalWithLocal” operation is used to update the global model with the changes of the local model. “syncLocalWithGlobal” operation is then used to download new local model data from the updated global table. At the end of the iteration, the sum of word counts for each topic is calculated with “allreduce” operation [11].

In “lda-rtt”, each worker will first conduct sampling with the global model partitions owned by itself and update them directly. Then it will call “rotateGlobal” operation to send the updated model data to the right neighbor and receive model

partitions from the left neighbor. Once all partitions of the global model are received and processed, the sampling of one iteration is completed. Similar to “lda-lgs”, “allreduce” operation is used at the end of the iteration to update the global sum of word counts on all topics.

C. Overlap Communication with Computation

Synchronous communication methods are often criticized for generating much overhead and making all workers wait for the completion of synchronization. We approached this problem in three steps. The first step is to balance the communication load on each worker through partitioning the global model based on word frequencies. The second step is to improve the speed with optimized collective communication. Here we discuss the third step, which is overlapping computation and communication in execution.

In “lda-rtt”, we slice the global model partitions held on each worker into two sets. Slicing is conducted by first sorting the partition IDs in ascending order and then assigning the partitions to the two slices in alternate order. As a result, each slice will contain words with both high and low frequencies. During the sampling, when a worker finishes processing the first slice, it uses another thread to rotate this slice and simultaneously continues processing the second slice. Once the second slice is processed, the first slice may be ready for further processing. Once both slices have finished a round of rotation, the sampling of an iteration is completed. The overlapping between computation and communication occurs when the worker processes one slice and rotates another slice at the same time.

In “lda-lgs”, we split the local data table into two slices. During the sampling, when each worker samples a slice, it requests another thread to synchronize the other slice through “syncLocalWithGlobal” and “syncGlobalWithLocal” operations. We map partitions based on their IDs into slices so that local partitions with the same IDs are guaranteed to be synchronized in iterations.

D. Inner-node Parallelism

In Harp-LDA, we use the “Computation” component to manage multi-threading sampling within one worker. The sampling process follows a SparseLDA algorithm and can be performed in two ways. One approach is to go through each document and sample the topics of every token. The other approach is to go through each word and sample the topics for word occurrences in each document. To keep the sampling order consistent between implementations for unbiased performance comparisons in future experiments, we sample topics by documents in “lda-lgs” as Yahoo! LDA and sample topics by words in “lda-rtt” as Petuum LDA. Note that when sampling topics by words, we balance the computation load by assigning words to threads based on their frequencies.

The local model is shared between threads. When sampling topics by documents, the word-topic model is required to access with locks. Symmetrically, when sampling topics by words, the document-topic model is required to access with

locks. We provide a read lock and a write lock on each document/word’s topic count map. Before sampling, a token’s document/word topic counts are read out, and after sampling, the updates are written back. If the next token for sampling is the same word, the sampling thread will keep using the thread local cached topic counts to avoid repeating fetching the shared data. During the update, we separate “updating an existing topic entry” and “adding a count to a new topic entry”. In “updating an existing topic entry”, because the map structure is not altered during updating and reading a primitive integer is an atomic operation in modern x86 architecture, it is safe to execute “read” and “update” concurrently with a shared read lock. However, in order to ensure the correctness of the topic count values, “update” operations are still required to be exclusive. In the operation of “adding a count to a new topic entry”, since the map structure is modified, we have to use a write lock.

Though the concurrency is greatly improved, our current implementation is still slower compared with Yahoo! LDA and Petuum in the first iteration of sampling. This could be caused by the difference in the implementation language (Java/C++) and the performance of the data structure (primitive int based hashmap [12]/primitive int array). As many-core architecture is becoming more common, high performance concurrent sampling with many-threads is a challenge to all implementations. However, in this paper our aim is not to provide the fastest LDA implementation but to show the advantages of using synchronous communication methods in LDA model convergence compared with asynchronous communication methods.

V. EXPERIMENTS

A. Experiment Settings

Experiments are done on the Juliet cluster [13], which contains 32 18-core 72-thread nodes and 96 24-core 48-thread nodes. All the nodes have 128GB memory and are connected with two types of networks: 1Gbps Ethernet (eth) and 16Gbps Infiniband (ib). For testing, we use 31 18-core nodes and 69 24-core nodes to form a cluster of 100 nodes with 40 threads each for computation. Most tests are done with Infiniband through IPoIB support unless otherwise specified.

Several datasets are used (see Table I). The total number of model parameters is kept as 10 billion on all datasets. α and β are both fixed at 0.01. We test several implementations (see Table II) on these datasets. We compare synchronous communication methods with asynchronous communication methods on both model parallelism and data parallelism. By studying the convergence speed and execution time, we learned how the difference in communication methods affects the performance of LDA.

B. Convergence Speed Per Iteration

First, we compare the convergence speed of the LDA word-topic model on iterations by analyzing model results learned on iterations 1, 10, 20, 30... 200. It is fair to use iterations to measure the performance of model convergence because it does not consider the performance difference in execution.

TABLE I
TRAINING DATA SETTINGS USED IN THE EXPERIMENTS

Dataset	enwiki	clueweb	bi-gram	gutenberg
Num. of Docs	3.8M	50.5M	3.9M	26.2K
Num. of Tokens	1.1B	12.4B	1.7B	836.8M
Vocabulary	1M	1M	20M	1M
Doc Len. AVG/STD	293/523	224/352	434/776	31879/42147
Lowest Word Freq.	7	285	6	2
Num. of Topics	10K	10K	500	10K
Init. Model Size	2.0GB	14.7GB	5.9GB	1.7GB

Note: Both “enwiki” and “bi-gram” are English articles from Wikipedia. “clueweb” is a 10% dataset from ClueWeb09, which is a collection of English web pages [14]. “gutenberg” is comprised of English books from Project GutenBurg [15].

TABLE II
LDA IMPLEMENTATIONS USED IN THE EXPERIMENTS

DATA PARALLELISM	
lgs	- “lda-lgs” impl. with no routing optimization - Slower than “lgs-opt”
lgs-opt	- “lgs” with routing optimization - Faster than Yahoo! LDA on “enwiki” with higher model likelihood
lgs-opt-4s	- “lgs-opt” with 4 rounds of model synchronization per iteration; each round uses 1/4 of the training data - Performance comparable to Yahoo! LDA on “clueweb” with higher model likelihood
Yahoo! LDA	- Master branch on GitHub [3]
MODEL PARALLELISM	
rtt	- “lda-rtt” impl. - Speed comparable with Petuum on “clueweb” but 4 times faster on “bi-gram” and “gutenberg”
Petuum	- Version 1.1 [4]

Note: Our implementations are indicated in bold.

On the “clueweb” dataset (see Fig. 4a), Petuum has the highest model likelihood on all iterations. Though “rtt” also uses model parallelism, due to its preference of using the thread-local data and not the up-to-date local shared model, the convergence speed is slower. “rtt” and “lgs-opt” have similar convergence speeds, and their lines on the chart are overlapped. In contrast to “lgs-opt”, the convergence speed of “lgs-opt-4s” is as high as Petuum. This shows that increasing the rounds of model synchronization thereby increases the convergence speed. Yahoo! LDA has the slowest convergence speed because asynchronous communication does not guarantee a full model synchronization in a iteration.

On the “enwiki” dataset (see Fig. 4b), as before, Petuum achieved the highest accuracy out of all iterations. “rtt” converges to the same model likelihood level as Petuum at iteration 200. “lgs-opt” demonstrates slower convergence speed but still achieved high model likelihood, while Yahoo! LDA has both the slowest convergence speed and lowest model likelihood at iteration 200.

All these results show that when the model update rate is increased (either using the model parallelism or using multiple-rounds model synchronization in data parallelism), the model converges faster.

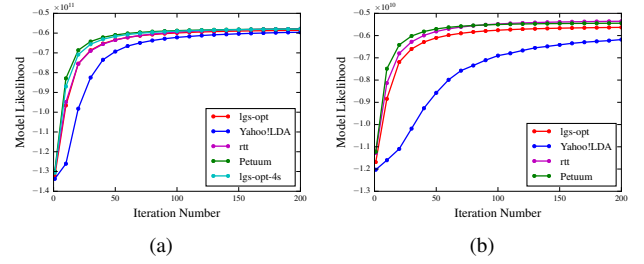


Fig. 4. Model Convergence of (a) “clueweb” And (b) “enwiki” On Iterations

C. Performance Analysis on Data Parallelism

We compare the model convergence speed on “lgs” and Yahoo! LDA by injecting the real execution time on iterations. On the “clueweb” dataset, we first show the convergence speed based on elapsed execution time (see Fig. 5a). Yahoo! LDA needs more time to obtain the model result of iteration 1 due to its slow model initialization. Since model initialization is mainly communication rather than computation and cannot be overlapped with sampling, Yahoo! LDA has a sizable overhead on the communication end. In later iterations, though “lgs” converges faster, Yahoo! LDA catches up after 30 iterations. This observation can be explained by our slower concurrent sampling speed and the fact that we only allow one round of model synchronization per iteration, while Yahoo! LDA does not have this restriction and allow multiple instances of synchronization whenever possible. Our computation takes quite long and the network is often in an idle state, therefore, we can increase the rounds of model synchronization per iteration. Although the execution time of 200 iterations for “lgs-opt-4s” is still slightly longer than Yahoo! LDA, it obtains higher model likelihood and maintains faster convergence speed in the whole execution.

Due to the slowness of the local concurrent sampling, our implementations show much higher iteration execution time at the first iteration compared with Yahoo! LDA (see Fig. 5b). However, with optimized synchronous communication methods, we quickly reduced the difference in execution time compared with Yahoo! LDA. Similar results are also shown on the “enwiki” dataset. “lgs-opt” not only achieves higher model likelihood but also has faster model convergence speed throughout the whole execution (see Fig. 5d). Though our execution time at iteration 1 is twice as slow as Yahoo! LDA, later on it takes less execution time per iteration than Yahoo! LDA (see Fig. 5e). Yahoo! LDA only exceeds “lgs-opt” when both models converge to a similar likelihood level.

In addition, we examined the effectiveness of using routing optimization in our “lgs” solution. Fig. 5c and Fig. 5f show that on Ethernet, “lgs-opt” is obviously faster than “lgs”; with Infiniband, due to its high bandwidth, the performance is very close to one another.

D. Performance Analysis on Model Parallelism

Here we compare “rtt” and Petuum on 3 different datasets: “clueweb”, “bi-gram”, and “gutenberg”. Since both implemen-

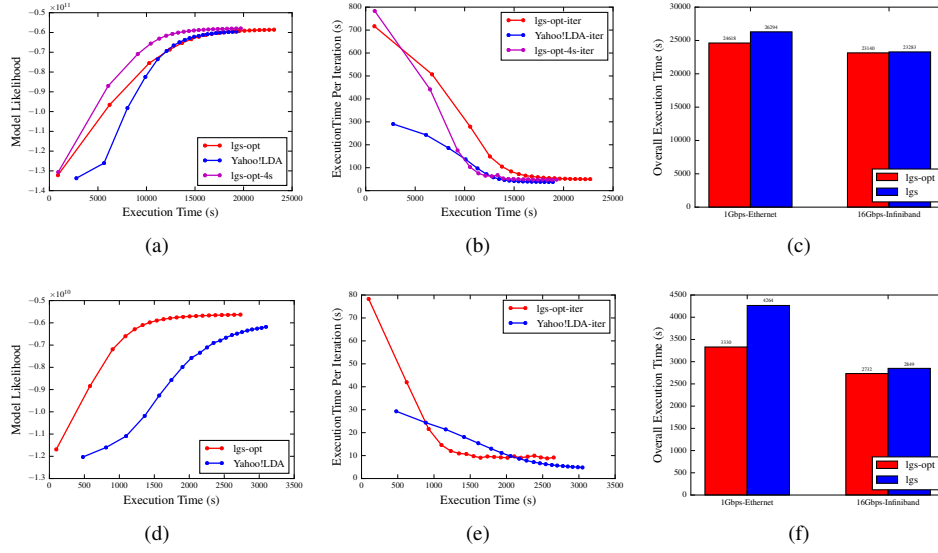


Fig. 5. Performance comparison on data parallelism between “lgs” and Yahoo! LDA (a) Elapsed Execution Time vs. Model Likelihood on “clueweb” (b) Elapsed Execution Time vs. Iteration Execution Time on “clueweb” (c) Total 200-Iteration Execution Time with Routing Optimization vs. without Routing Optimization on “clueweb” with ib/eth (d) Elapsed Execution Time vs. Model Likelihood on “enwiki” (e) Elapsed Execution Time vs. Iteration Execution Time on “enwiki” (f) Total 200-Iteration Execution Time with Routing Optimization vs. without Routing Optimization on “enwiki” with ib/eth

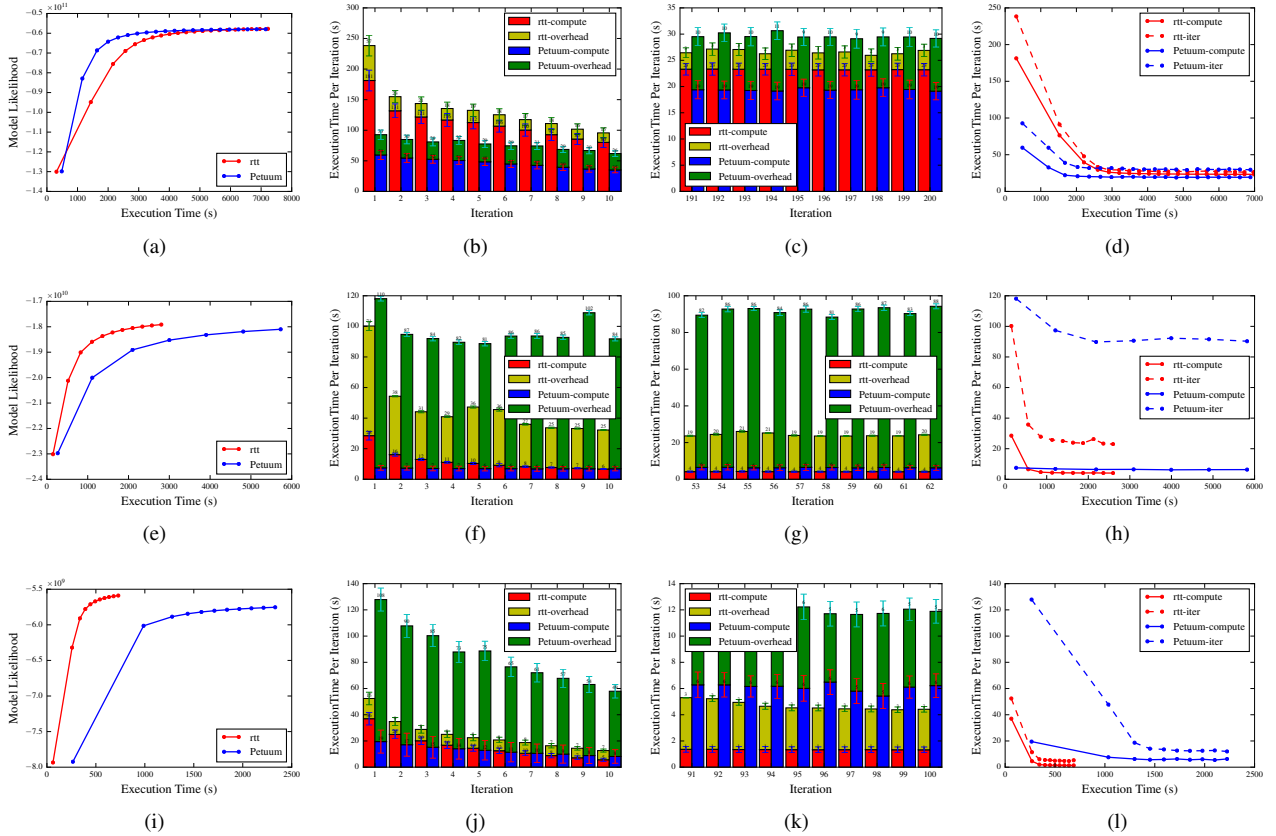


Fig. 6. Performance comparison on model parallelism between “rtt” and Petuum (a) Elapsed Execution Time vs. Model Likelihood on “clueweb” (b) First 10 Iteration Execution Times on “clueweb” (c) Final 10 Iteration Execution Times on “clueweb” (d) Elapsed Execution Time vs. Iteration Execution Time on “clueweb” (e) Elapsed Execution Time vs. Model Likelihood on “bi-gram” (f) First 10 Iteration Execution Times on “bi-gram” (g) Final 10 Iteration Execution Times on “bi-gram” (h) Elapsed Execution Time vs. Iteration Execution Time on “bi-gram” (i) Elapsed Execution Time vs. Model Likelihood on “gutenberg” (j) First 10 Iteration Execution Times on “gutenberg” (k) Final 10 Iteration Execution Times on “gutenberg” (l) Elapsed Execution Time vs. Iteration Execution Time on “gutenberg”

tations use model parallelism, the performance difference is caused by the execution speed per iteration.

On the “clueweb” dataset, the execution times after 200 iterations were similar between both implementations, and they achieved similar model likelihood (see Fig. 6a). The first 10 iterations show that “rtt” has high computation time compared with Petuum (see Fig. 6b), however, its overhead on communication per iteration becomes lower than Petuum. When the execution arrives at the final 10 iterations, while computation overhead per iteration in “rtt” is still higher, the whole execution time per iteration becomes lower (see Fig. 6c). The trend of the iteration execution time on 200 iterations is shown in Fig 6d.

Unlike our “rotateGlobal” operation which batches transmission of model data partitions, Petuum sends model data word by word asynchronously, causing high communication overhead. On the “bi-gram” dataset, the results show that Petuum cannot perform well when the number of words in the model increases. The high overhead in communication causes the convergence speed to be very slow, and Petuum cannot even continue executing after 60 iterations due to a memory outage (see Fig. 6e). Fig. 6f and Fig. 6g show that in the first/final 10 iterations, Petuum consistently has higher execution time per iteration compared with “rtt”. The trend of the iteration execution time on 200 iterations also shows this phenomenon (see Fig. 6h).

Though the data size of “gutenberg” is similar to “enwiki”, it is clear that there is a difference in execution speed per iteration (see Fig. 6i). High standard deviation indicates that the iteration execution time per worker varies significantly. Unlike the results on “bi-gram” where Petuum’s performance suffers from the communication overhead, here it suffers from waiting for the slowest worker. The high iteration execution time may be explained by “gutenberg” containing many long documents and thereby resulting in unbalanced training data distribution on the workers. In addition, when sampling by words, frequent access to the shared huge doc-topic model leads to inefficient concurrent sampling. “rtt” is not much affected because it prefers using thread-local data in concurrent sampling and balances per-thread computation through assigning words to threads based on the frequencies. Fig. 6j, Fig. 6k, and Fig. 6l display that the unbalanced computation in Petuum results in high overhead per iteration. In model parallelism, model rotation is a synchronous operation; therefore, this experiment demonstrates that unbalanced computation on workers causes huge overhead in global waiting and results in high iteration execution time. In sum, when applying synchronous communication methods, the computation load should be carefully balanced.

VI. RELATED WORK

Prior research has studied the parallelization of the LDA algorithm extensively. Some studies focused on using the Collapsed Variational Bayes (CVB) algorithm [1]. Mahout LDA [16] and Spark LDA [17] both use this algorithm.

TABLE III
LDA WORK USING CGS ALGORITHM

App. Name	Algorithm	Parallelism	Comm.
PLDA	CGS (sample by docs)	D. P.	allreduce (sync)
Dato	CGS (sample by doc-word edge)	D. P.	GAS (sync)
Yahoo! LDA	CGS (SparseLDA & sample by docs)	D. P.	client-server (async)
Peacock	CGS (SparseLDA & sample by words)	D. P. (M. P. in local)	client-server (async)
Parameter Server	CGS (combined with other methods)	D. P.	client-server (async)
Petuum 0.93	CGS (SparseLDA & sample by docs)	D. P.	client-server (async)
Petuum 1.1	CGS (SparseLDA & sample by words)	M. P. (include D. P.)	ring/star topology (async)

Note: “D. P.” refers to Data Parallelism. “M. P.” refers to Model Parallelism.

However, research also shows that this approach leads to high memory consumption and slow convergence speed [6][18].

Other studies use the CGS algorithm (see Table III). PLDA [19] is such an implementation. There are two versions of PLDA, one based on MPI [20] using the “allreduce” operation [11], and the other based on on MapReduce[10][21] using “shuffle” operation.

Yahoo! LDA [22][23] uses the CGS algorithm with SparseLDA optimization, and its architecture is client-server with asynchronous communication. Local models are distributed in the star model, and local computation threads use optimized locking mechanisms when accessing the shared local model. The synchronization between local models and the global model is done through asynchronous delta aggregation.

Dato [24] uses the GAS model [25] to implement the LDA algorithm [26]. Currently, it uses a CGS algorithm without SparseLDA optimization. GAS model’s edge-based computation patterns cause the training data to be partitioned based on document-word pairs instead of the documents. As a result, during the sampling process, both the topic counts of words and documents have to be gathered and updated. This results in additional communication costs in synchronization.

Peacock [18] uses a hierarchical distributed architecture to organize the LDA computation. The first layer uses the SparseLDA algorithm with a lock-free parallel strategy to exploit local model parallelism. The design of this layer is similar to “rotateGlobal” but differs by sending documents to where the model locates rather than rotating model partitions between documents. The second layer also uses client-server architecture with asynchronous communication.

Parameter Server [27] and Petuum [28] both provide a framework to allow programming machine learning algorithms in client-server architecture with “push” and “pull” operations. Parameter Server puts the global model on servers and uses range-based “push” and “pull” operations for synchronization.

These operations allow workers to update a row or a segment of parameters directly and provides a chance to batch the communication of model updates. The computation of Parameter Server's LDA implementation uses a combination of stochastic variational methods, collapsed Gibbs sampling, and distributed gradient descent. Another operation of Petuum, "schedule", allows model parallelism through scheduling model partitions to workers. Lee et al. [29] describes that the communication to fetch model data goes between clients and servers, but in the real code on GitHub [4], workers are actually directly sending data to neighbors with optimized routing.

VII. CONCLUSION

Through experiments on several datasets, we showed that synchronous communication methods perform better than asynchronous methods on both data parallelism and model parallelism. In data parallelism, our implementation with synchronous communication methods resulted in faster model convergence and higher model likelihood at the final iteration compared to Yahoo! LDA using asynchronous communication methods. In model parallelism, our implementation with synchronous communication methods also showed significantly lower overhead than Petuum LDA. On "bi-gram" dataset, the total execution time of "rtt" is four times faster than Petuum. Even though the computation speed of the first iteration is 2- to 3-fold slower on "clueweb" dataset, the total execution time of "rtt" remains similar with Petuum. These results prove that with optimized synchronous communication methods, we can increase the model update rate, allowing the model to converge faster, shrinking the model size, and further reducing the computation time in later iterations.

In general, despite the implementation differences in performance between "rtt", "lgs", Yahoo! LDA, and Petuum LDA, the advantages of synchronous communication methods are obvious. Compared with asynchronous communication methods, synchronous communication methods can optimize routing between a set of parallel workers and maximize bandwidth utilization in point-to-point communication. Synchronous communication methods may result in global/local waiting. However, since the word frequencies in the LDA training data is under the power-law distribution and a considerable amount of words have high frequencies, balancing the computation on all parallel workers is feasible, and the overhead of waiting is not as high as speculated. The chain reaction set off by improving the LDA model update speed amplifies the benefit of using synchronous communication methods.

In future work, we will focus on improving inner-node model synchronization speed in many-core systems to provide a high performance LDA implementation and apply our model synchronization strategies to other machine learning algorithms facing difficulties in handling big model data.

ACKNOWLEDGMENT

We appreciate the system support offered by FutureSystems. We gratefully acknowledge support from NSF 1443054 SPI-

DAL grant.

REFERENCES

- [1] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *The Journal of Machine Learning Research*, vol. 3, pp. 993–102, 2003.
- [2] B. Zhang, Y. Ruan, and J. Qiu, "Harp: collective communication on hadoop," in *IC2E*, 2015.
- [3] "Yahoo! LDA." [Online]. Available: https://github.com/sudar/Yahoo_LDA
- [4] "Petuum LDA." [Online]. Available: <https://github.com/petuum/bosen/wiki/Latent-Dirichlet-Allocation>
- [5] P. Resnik and E. Hardist, "Gibbs sampling for the uninitiated," University of Maryland, Tech. Rep., 2010.
- [6] D. Newman, A. Asuncion, P. Smyth, and M. Welling, "Distributed algorithms for topic models," *The Journal of Machine Learning Research*, vol. 10, pp. 1801–1828, 2009.
- [7] J. Yuan, F. Gao, Q. Ho, W. Dai, J. Wei, X. Zheng, E. P. Xing, T.-Y. Liu, and W.-Y. Ma, "LightLDA: big topic models on modest computer clusters," in *WWW*, 2015, pp. 1351–1361.
- [8] L. Yao, D. Mimno, and A. McCallum, "Efficient methods for topic model inference on streaming document collections," in *KDD*, 2009, pp. 937–946.
- [9] C.-T. Chu, S. K. Kim, Y.-A. Lin, Y. Yu, G. Bradski, A. Y. Ng, and K. Olukotun, "Map-reduce for machine learning on multicore," in *NIPS*, vol. 19, 2007, p. 281.
- [10] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [11] E. Chan, M. Heimlich, A. Purkayastha, and R. V. D. Geijn, "Collective communication: theory, practice, and experience," *Concurrency and Computation: Practice and Experience*, vol. 19, no. 13, pp. 1749–1783, 2007.
- [12] "fastutil." [Online]. Available: <http://fastutil.di.unimi.it>
- [13] "FutureSystems." [Online]. Available: <https://portal.futuresystems.org>
- [14] "clueweb." [Online]. Available: <http://boston.lti.cs.cmu.edu/clueweb09/wiki/tiki-index.php?page=Dataset+Information>
- [15] "gutenberg." [Online]. Available: <https://www.gutenberg.org>
- [16] "Mahout LDA." [Online]. Available: <https://mahout.apache.org/users/clustering/latent-dirichlet-allocation.html>
- [17] "Spark LDA." [Online]. Available: <http://spark.apache.org/docs/latest/ml-lib-clustering.html>
- [18] Y. Wang, X. Zhao, Z. Sun, H. Yan, L. Wang, Z. Jin, L. Wang, Y. Gao, C. Law, and J. Zeng, "Peacock: learning long-tail topic features for industrial applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 6, no. 4, 2015.
- [19] Y. Wang, H. Bai, M. Stanton, W.-Y. Chen, and E. Y. Chang, "PLDA: parallel latent dirichlet allocation for large-scale applications," *Algorithmic Aspects in Information and Management*, pp. 301–314, 2009.
- [20] D. W. Walker and J. J. Dongarra, "MPI: a standard message passing interface," in *Supercomputer*, vol. 12, 1996, pp. 56–68.
- [21] "Hadoop." [Online]. Available: <http://hadoop.apache.org>
- [22] A. Smola and S. Narayanamurthy, "An architecture for parallel topic models," in *VLDB*, vol. 3, no. 1-2, 2010, pp. 703–710.
- [23] A. Ahmed, M. Aly, J. Gonzalez, S. Narayanamurthy, and A. J. Smola, "Scalable inference in latent variable models," in *WSDM*, 2012, pp. 123–132.
- [24] "Dato." [Online]. Available: <https://dato.com>
- [25] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin, "PowerGraph: distributed graph-parallel computation on natural graphs," in *OSDI*, vol. 12, 2012, p. 2.
- [26] "Dato LDA." [Online]. Available: https://github.com/dato-code/PowerGraph/blob/master/toolkits/topic_modeling/topic_modeling.dox
- [27] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su, "Scaling distributed machine learning with the parameter server," in *OSDI*, 2014, pp. 583–598.
- [28] E. P. Xing, Q. Ho, W. Dai, J. K. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Yu, "Petuum: a new platform for distributed machine learning on big data," in *KDD*, 2013.
- [29] S. Lee, J. K. Kim, X. Zheng, Q. Ho, G. A. Gibson, and E. P. Xing, "On model parallelization and scheduling strategies for distributed machine learning," in *NIPS*, 2014, pp. 2834–2842.